Information Diffusion in Social Sensing

Vikram Krishnamurthy, *Fellow, IEEE*, William Hoiles, *Student Member, IEEE*

arXiv:1412.4171v1 [cs.SI] 13 Dec 2014

***Abstract*—Statistical inference using social sensors is an area that has witnessed remarkable progress in the last decade. It is relevant in a variety of applications including localizing events for targeted advertising, marketing, localization of natural disasters and predicting sentiment of investors in financial markets. This paper presents a tutorial description of three important aspects of sensing-based information diffusion in social networks from a communications/signal processing perspective. First, diffusion models for information exchange in large scale social networks together with social sensing via social media networks such as Twitter is considered. Second, Bayesian social learning models in online reputation systems are presented. Finally, the principle of revealed preferences arising in micro-economics theory is used to parse datasets to determine if social sensors are utility maximizers and then determine their utility functions. All three topics are explained in the context of actual experimental datasets from health networks, social media and psychological experiments. Also, algorithms are given that exploit the above models to infer underlying events based on social sensing. The overview, insights, models and algorithms presented in this paper stem from recent developments in computer-science, economics, psychology and electrical engineering.**

Keywords: diffusion, Susceptible-Infected-Susceptible (SIS model), random graphs, mean field dynamics, social sampling, social learning, reputation systems, Afriat's theorem, detecting utility maximizers, revealed preferences, potential games

I. Introduction and Motivation

Humans can be viewed as social sensors that interact over a social network to provide information about their environment. Examples of information produced by such social sensors include Twitter posts, Facebook status updates, and ratings on online reputation systems like Yelp and Tripadvisor. Social sensors go beyond physical sensors – for example, user opinions/ratings (such as the quality of a restaurant) are available on Tripadvisor but are difficult to measure via physical sensors. Similarly, future situations revealed by the Facebook status of a user are impossible to predict using physical sensors [1].

Statistical inference using social sensors is an area that has witnessed remarkable progress in the last decade. It is relevant in a variety of applications including localizing special events for targeted advertising [2], [3], marketing [4], [5], localization of natural disasters [6], and predicting sentiment of investors in financial markets [7], [8]. For example, [9] reports sthat models built from the rate of tweets for particular products can outperform market-based predictors.

A. Context: Why social sensors?

Social sensors present unique challenges from a statistical estimation point of view. First, social sensors interact with and influence other social sensors. For example, ratings posted on online reputation systems strongly influence the behavior of individuals.[1] Such interacting sensing can result in non-standard information patterns due to correlations introduced by the structure of the underlying social network. Thus certain events "go viral" [5], [12]. Second, due to privacy concerns and time-constraints, social sensors typically do not reveal raw observations of the underlying state of nature. Instead, they reveal their decisions (ratings, recommendations, votes) which can be viewed as a low resolution (quantized) function of their raw measurements and interactions with other social sensors. This can result in misinformation propagation, herding and information cascades. Third, the response of a social sensors may not be consistent with that of an utility maximizer.

Social sensors are enabled by technological networks. Indeed, social media sites that support interpersonal communication and collaboration using Internet-based social network platforms, are growing rapidly. The following excerpt from [13] illustrates the increasing importance of social sensors in marketing:

- "53% of people on Twitter recommend companies and/or products in their tweets, with 48% of them delivering on their intention to buy the product. (ROI Research for Performance, June 2010)

V. Krishnamurthy (e-mail: vikramk@ece.ubc.ca) and W. Hoiles (e-mail whoiles@eceubc.ca) are with the Department of Electrical and Computer Engineering, University of British Columbia, Vancouver, V6T 1Z4, Canada.

This research was supported by the Canada Research Chairs program, Natural Sciences and Engineering Research Council of Canada and Social Sciences and Humanities Research Council of Canada.

[1] It is reported in [10] that 81% of hotel managers regularly check Tripadvisor reviews. [11] reports that a one-star increase in the Yelp rating maps to 5-9 % revenue increase.

- The average consumer mentions specific brands over 90 times per week in conversations with friends, family, and co-workers. (Keller Fay, WOMMA, 2010)
- Consumer reviews are significantly more trusted—nearly 12 times more—than descriptions that come from manufacturers, according to a survey of US mom Internet users by online video review site EXPO. (eMarketer, February 2010)
- In a study conducted by social networking site myYearbook, 81% of respondents said they had received advice from friends and followers relating to a product purchase through a social site; 74% of those who received such advice found it to be influential in their decision. (ClickZ, January 2010)

McKinsey estimates that the economic impact of social media on business is potentially greater than $1 trillion since social media facilitates efficient communication and collaboration within and across organizations.

B. Main Results and Organization

There is strong motivation to construct models that facilitate understanding the dynamics of information flow in social networks.This paper presents a tutorial description of three important aspects of sensing-based information diffusion in social networks from a communications/signal processing perspective:

1) Information Diffusion in Large Scale Social Networks: The first topic considered in this paper (Sec.II) is diffusion of information in social networks comprised of a population of interacting social sensors. The states of sensors evolve over time as a probabilistic function of the states of their neighbors and an underlying target process. Several recent papers investigate such information diffusion in real-world social networks. Motivated by marketing applications, [14] studies diffusion (contagion) behavior in Facebook. Using data from 260,000 Facebook *pages* (which advertise products, services and celebrities), [14] analyzes information diffusion. In [15], the spread of *hashtags* on Twitter is studied. There is a wide range of social phenomena such as diffusion of technological innovations, sentiment, cultural fads, and economic conventions [16], [17] where individual decisions are influenced by the decisions of others.

We consider the so called Susceptible-Infected-Susceptible (SIS) model [18] for information diffusion in a social network. It is shown for social networks comprised of a large number of agents how the dynamics of degree distribution can be approximated by the mean field dynamics. Mean field dynamics have been studied in [19] and applied to social networks in [17] and leads to a tractable model for the dynamics social sensors.

We demonstrate using influenza datasets from the U.S Centers for Disease Control and Prevention (CDC) how Twitter can be used as a real time social sensor for tracking the spread of influenza. That is, a health network (namely, Influenza-like Illness Surveillance Network (ILInet)) is sensed by a real time microblogging social media network (namely, Twitter).

We also review two recent methods for sampling social networks, namely, social sampling and respondent-driven sampling. Respondent-driven sampling is now used by the U.S. Centers for Disease Control and Prevention (CDC) as part of the National HIV Behavioral Surveillance System in health networks.

2) Bayesian Social Learning in Online Reputation Systems: The second topic of this paper (Sec.III) considers online reputation systems where individuals make recommendations based on their private observations and recommendations of friends. Such interaction of individuals and their social influence is modelled as Bayesian social learning [20], [21], [16] on a directed acyclic graph. Data incest (misinformation propagation) arises as a result of correlations in recommendations due to the intersection of multiple paths in the information exchange graph. Necessary and sufficient conditions are given on the structure of information exchange graph to mitigate data incest. Experimental results on human subjects are presented to illustrate the effect of social influence and data incest on decision making.

The setup differs from classical signal processing where sensors use noisy observations to compute estimates - in social learning agents use noisy observations together with decisions made by previous agents, to estimate the underlying state of nature. Social learning has been used widely in economics, marketing, political science and sociology to model the behavior of financial markets, crowds, social groups and social networks; see [20], [21], [22], [16], [23], [24] and numerous references therein. Related models have been studied in the context of sequential decision making in information theory [25], [26] and statistical signal processing [27], [28] in the electrical engineering literature. Social learning can result in unusual behavior such as herding [21] where agents eventually choosing the same action irrespective of their private observations. As a result, the actions contain no information about the private observations and so the Bayesian estimate of the underlying random variable freezes. Such behavior can be undesirable, particularly if individuals herd and make incorrect decisions.

3) Revealed Preferences and Detection of Utility Maximizers: The third topic considered in this paper

(Sec.IV) is the principle of revealed preferences arising in microeconomics. It is used as a constructive test to determine: Are social sensors utility optimizers in their response to external influence? The key question considered is as follows: Given a time-series of data $\mathcal{D} = \{(p_t, x_t), t \in \{1, 2, \ldots, T\}\}$ where $p_t \in \mathbb{R}^m$ denotes the external influence, x_t denotes the response of an agent, is it possible to detect if the agent is a *utility maximizer*?

These issues are fundamentally different to the *model-centric* theme used in the telecommunications literature where one postulates an objective function (typically convex) and then proposes optimization algorithms. In contrast the revealed preference approach is *data centric* - given a dataset, we wish to determine if is consistent with utility maximization.

We present a remarkable result called Afriat's theorem [29], [30] which provides a necessary and sufficient condition for a finite dataset $\mathcal{D}$ to have originated from a utility maximizer. Also a multi-agent version of Afriat's theorem is presented to determine if the dataset generated by multiple agents is consistent with playing from the equilibrium of a potential game.

Unlike model centric applications of game theory in telecommunications, our revealed preferences approach is data centric: 1) Given a time series dataset of probe and response signals, how can one detect if the response signals are consistent with a Nash equilibrium generated by players in a concave potential game? 2) If consistent with a concave potential game, how can the utility function of the players be estimated?

We present three datasets involving social sensors to illustrate Afriat's theorem of revealed preferences. These datasets are: (i) an auction conducted by undergraduate students at Princeton University, (ii) aggregate power consumption in the electricity market of Ontario province and (iii) Twitter dataset for specific hashtags.

Varian (chief economist at Google) has written several influential papers on Afriat's theorem in the economics literature. These include measuring the welfare effect of price discrimination [31], analysing the relationship between prices of broadband Internet access and time of use service [32], and ad auctions for advertisement position placement on page search results from Google [32], [33]. Despite widespread use in economics, revealed preference theory is relatively unknown in the electrical engineering literature.

C. Perspective

The unifying theme that underpins the three topics in this paper stems from predicting global behavior given local behavior: individual social sensors interact with other sensors and we are interested in understanding the behavior of the entire network. Information diffusion, social learning and revealed preferences are important issues for social sensors. We treat these issues in a highly stylized manner so as to provide easy accessibility to a communications/signal processing audience. The underlying tools used in this paper are widely used in signal processing, control, information theory and network communications.

Social Sensors vs Cognitive Radio: There are interesting parallels between social sensors (considered in this paper) and spectrum sensing in cognitive radio. Cognitive radios [34] scan spectrum, determine spectral gaps, and then collaborate with other cognitive radios to determine the best spectrum gap to transmit in. A social sensor uses cognition to choose the most relevant content in a social media; that is, it scans social media (termed diversity spectrum in [35]), then interacts with other users to determine the most relevant information, and then broadcasts it in the social media. Of course, with social sensors there is the additional complexity that the information communicated diffuses and influences other social sensors to re-post that content; this can result in a cascading behavior [36].

Books and Tutorials: The literature in social learning, information diffusion and revealed preferences is extensive. In each of the following sections, we provide a brief review of relevant works. Seminal books in social networks, social learning and network science include [37], [38], [16], [39]. There is a growing literature dealing with the interplay of technological and social networks [40]. Social networks overlaid on technological networks account for a significant fraction of Internet use. As discussed in [40], three key aspects that cut across social and technological networks are the emergence of global coordination through local actions, resource sharing models and the wisdom of crowds. These themes are addressed in the current paper in the context of social learning, diffusion and revealed preferences. Other tutorials include [41], [42].

II. Information Diffusion in Large Scale Social Networks

This section addresses the first topic of the paper, namely, information diffusion models and their mean field dynamics in social networks. The setting is as follows: The states of individual nodes in the social network evolve over time as a probabilistic function of the states of their neighbors *and* an underlying target process (state of nature). The underlying target process

can be viewed as the market conditions or competing technologies that evolve with time and affect the information diffusion. The nodes in the social network are sampled randomly to determine their state. As the adoption of the new technology diffuses through the network, its effect is observed via sentiments (such as tweets) of these selected members of the population. These selected nodes act as social sensors. In signal processing terms, the underlying target process can be viewed as a signal, and the social network can be viewed as a sensor. The key difference compared to classical sensing is that the sensor now is a social network with diffusion dynamics and noisy measurements (due to sampling nodes).

As described in Sec.I, a wide range of social phenomena such as diffusion of technological innovations, cultural fads, ideas, behaviors, trends and economic conventions [43], [44], [45], [16] can be modelled by diffusion in social networks. Another important application is sentiment analysis (opinion mining) where the spread of opinions amongst people is monitored via social media.

Motivated by the above setting, this section proceeds as follows:

1) We describe the Susceptible-Infected-Susceptible (SIS) model for diffusion of information in social networks which has been extensively studied in [38], [46], [17], [18], [37].
2) Next, it is shown how the dynamics of the infected degree distribution of the social network can be approximated by the mean field dynamics. The mean field dynamics state that as the number of agents in the social network goes to infinity, the dynamics of the infected degree distribution converges to that of an ordinary differential (or difference) equation. Such averaging theory results are widely used to analyze adaptive filters. For social networks, they yield a useful tractable model for the diffusion dynamics.
3) We illustrate the diffusion model by using data sets from 3 related social networks to track the spread of influenza during the period September 1 to December 31, 2009. The friendship network of 744 undergraduate students at Harvard college is used together with the U.S. outpatient Influenza-like Illness Surveillance Network (ILInet) to monitor the spread of influenza. Then it is shown that Twitter posts related to influenza during this period are correlated with the spread of influenza. Thus in this example, influenza diffuses in a human health network (Harvard friendship network at a local level and ILInet at a global level) and Twitter is used as a social sensor to monitor the spread of the influenza.
4) Finally, this section also describes how social networks can be sampled. We review two recent methods for sampling social networks, namely, social sampling and respondent-driven sampling; the latter being used in health networks.

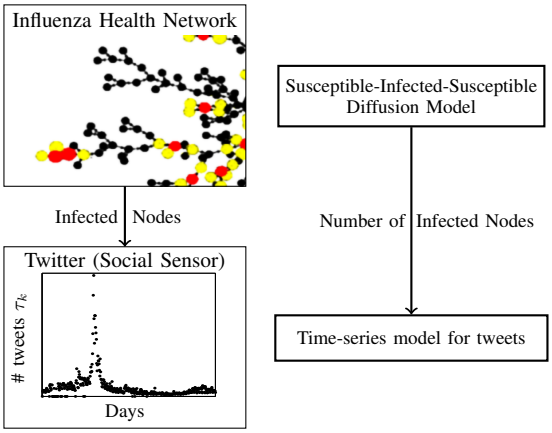

(a) Dynamics of Health Network impact the # of tweets. (b) Model for the number of tweets resulting from the Influenza Health Network.

Fig. 1. Dynamics of Health Network are modelled using the SIS model and a Linear and Nonlinear Autoregressive with Exogenous input time series models, refer to Sec.II for details.

The aim is to estimate the underlying target state that is being sensed by the social network and also and the state probabilities of the nodes by sampling measurements at nodes in the social network. In a Bayesian estimation context, this is equivalent to a filtering problem involving estimation of the state of a prohibitively large scale Markov chain in noise. The mean field dynamics yields a tractable approximation with provable bounds for the information diffusion. Such mean field dynamics have been studied in [19] and applied to social networks in [46], [17], [37]. For an excellent recent exposition of interacting particle systems comprising of agents each with a finite state space, see [47], where the more apt term "Finite Markov Information Exchange (FMIE) process" is used.

Regarding real datasets, in addition to the case study presented below, for other examples of diffusion datasets and their analysis see [14], [15]. A repository of social network datasets can be obtained at [48].

A. Social Network Model

A social network is modelled as a graph with N vertices:

$$G = (V, E), \text{ where } V = \{1, 2, \ldots, N\}, \text{ and } E \subseteq V \times V. \tag{1}$$

Here, V denotes the finite set of vertices, and E denotes the set of edges. In social networks, it is customary to use the terminology *network*, *nodes* and *links* for *graph*, *vertices* and *edges*, respectively.

We use the notation (m, n) to refer to a link between node m and n. The network may be undirected in which case $(m, n) \in E$ implies $(n, m) \in E$. In undirected graphs, to simplify notation, we use the notation m, n to denote the undirected link between node n and m. If the graph is directed, then $(m, n) \in E$ does not imply that $(n, m) \in E$. We will assume that self loops (reflexive links) of the form i, i are excluded from E.

An important parameter of a social network $G = (V, E)$ is the connectivity of its nodes. Let $\mathcal{N}^{(m)}$ and $D^{(m)}$ denote the neighbourhood set and degree (or connectivity) of a node $m \in V$, respectively. That is, with $|\cdot|$ denoting cardinality,

$$\mathcal{N}^{(m)} = \{n \in V : m, n \in E\}, \quad D^{(m)} = \left|\mathcal{N}^{(m)}\right|. \tag{2}$$

For convenience, we assume that the maximum degree of the network is uniformly bounded by some fixed integer $\bar{D}$.

Let $N(d)$ denote the number of nodes with degree d, and let the degree distribution $P(d)$ specify the fraction of nodes with degree d. That is, for $d = 0, 1, \ldots, \bar{D}$,

$$N(d) = \sum_{m \in V} I\left\{D^{(m)} = d\right\}, \quad P(d) = \frac{N(d)}{N}.$$

Here, $I\{\cdot\}$ denotes the indicator function. Note that $\sum_d P(d) = 1$. The degree distribution can be viewed as the probability that a node selected randomly with uniform distribution on V has a connectivity d.

Random graphs generated to have a degree distribution P that is Poisson were formulated by Erdös and Renyi [49]. Several recent works show that large scale social networks are characterized by connectivity distributions that are different to Poisson distributions. For example, the internet, www have a power law connectivity distribution $P(d) \propto d^{-\gamma}$, where γ ranges between 2 and 3. Such scale free networks are studied in [50]. In the rest of this chapter, we assume that the degree distribution of the social network is arbitrary but known—allowing an arbitrary degree distribution facilities modelling complex networks.

Let $k = 0, 1, \ldots$ denote discrete time. Assume the target process s is a finite state Markov chain with transition probability

$$A_{ss'} = \mathbb{P}\left(s_{k+1} = s' | s_k = s\right). \tag{3}$$

In the example of technology diffusion, the target process can denote the availability of competition or market forces that determine whether a node adopts the technology. In the model below, the target state will affect the probability that an agent adopts the new technology.

B. SIS Diffusion Model for Information in Social Network

The model we present below for the diffusion of information in the social network is called the *Susceptible-Infected-Susceptible (SIS)* model [18], [37]. The diffusion of information is modelled by the time evolution of the state of individual nodes in the network. Let $x_k^{(m)} \in \{0, 1\}$ denote the state at time k of each node m in the social network. Here, $x_k^{(m)} = 0$ if the agent at time k is susceptible and $x_k^{(m)} = 1$ if the agent is infected. At time k, the state vector of the N nodes is

$$x_k = \left[x_k^{(1)}, \ldots, x_k^{(N)}\right]' \in \{0, 1\}^N. \tag{4}$$

Assume that the process x evolves as a discrete time Markov process with transition law depending on the target state s. If node m has degree $D^{(m)} = d$, then the probability of node m switching from state i to j is

$$\begin{aligned} &\mathbb{P}\left(x_{k+1}^{(m)} = j | x_k^{(m)} = i, x_k^{(i-)}, s_k = s\right) \\ &\quad = p_{ij}(d, A_k^{(m)}, s),\ i, j \in \{0, 1\}. \end{aligned} \tag{5}$$

Here, $A_k^{(m)}$ denotes the number of infected neighbors of node m at time k. That is,

$$A_k^{(m)} = \sum_{n \in N^{(m)}} I\left\{n : x_k^{(m)} = 1\right\}. \tag{6}$$

In words, the transition probability of an agent depends on its degree distribution and the number of active neighbors.

With the above probabilistic model, we are interested in modelling the evolution of infected agents over time. Let $\rho_k(d)$ denote the fraction of infected nodes at each time k with degree d. We call ρ the *infected node distribution*. So with $d = 0, 1, \ldots, \bar{D}$,

$$\rho_k(d) = \frac{1}{N(d)} \sum_{m \in V} I\left\{D^{(m)} = d, x_k^{(m)} = 1\right\}. \tag{7}$$

The SIS model assumes that the infection spreads according to the following dynamics:

1) At each time instant k, a single agent, denoted by m, amongst the N agents is chosen uniformly. Therefore, the probability that the chosen agent m is infected and of degree d is $\rho_k(d)\,P(d)$. The probability that the chosen agent m is susceptible and of degree d is $(1-\rho_k(d))\,P(d)$.
2) Depending on whether its state $x_k^{(m)}$ is infected or susceptible, the state of agent m evolves according to the transition probabilities specified in (5).

With the Markov chain transition dynamics of individual agents specified above, it is clear that the infected distribution $\rho_k = \big(\rho_k(1),\ldots,\rho_k(\bar{D})\big)$ is an $\prod_{d=1}^{\bar{D}} N(d)$ state Markov chain. Indeed, given $\rho_k(d)$, due to the infection dynamics specified above

$$\rho_{k+1}(d) \in \left\{ \rho_k(d) - \frac{1}{N(d)},\ \rho_k(d) + \frac{1}{N(d)} \right\}. \quad (8)$$

Our aim below is to specify the transition probabilities of the Markov chain ρ. Let us start with the following statistic that forms a convenient parametrization of the transition probabilities. Given the infected node distribution ρ_k at time k, define $\theta(\rho_k)$ as the probability that at time k a uniformly sampled link in the network points to an infected node. We call $\theta(\rho_k)$ as the *infected link probability*. Clearly

$$\begin{aligned}\theta(\rho_k) &= \frac{\sum_{d=1}^{\bar{D}} (\text{\# of links from infected node of degree } d)}{\sum_{d=1}^{\bar{D}} (\text{\# of links of degree } d)} \\ &= \frac{\sum_{d=1}^{\bar{D}} d\,P(d)\,\rho_k(d)}{\sum_d^{\bar{D}} d\,P(d)}. \end{aligned} \quad (9)$$

In terms of the infected link probabilities, the scaled transition probabilities[2] of the process ρ are:

$$\begin{aligned}\bar{p}_{01}(d,\theta_k,s) &\stackrel{\text{defn}}{=} \frac{1}{P(d)}\,\mathbb{P}\left(\rho_{k+1}(d) = \rho_k(d) + \frac{1}{N(d)}\Big| s_k = s\right) \\ &= (1-\rho_k(d)) \sum_{a=0}^{d} p_{01}(d,a,s)\,\mathbb{P}(a \text{ out of } l \text{ neighbours infected}) \\ &= (1-\rho_k(d)) \sum_{a=0}^{d} p_{01}(d,a,s) \binom{d}{a} \theta_k^a (1-\theta_k)^{d-a}, \\ \bar{p}_{10}(d,\theta_k,s) &\stackrel{\text{defn}}{=} \frac{1}{P(d)}\,\mathbb{P}\left(\rho_{k+1}(d) = \rho_k(d) - \frac{1}{N(d)}\Big| s_k = s\right) \\ &= \rho_k(d) \sum_{a=0}^{d} p_{10}(d,a,s) \binom{d}{a} \theta_k^a (1-\theta_k)^{d-a}. \end{aligned} \quad (10)$$

In the above, the notation θ_k is the short form for $\theta(\rho_k)$. The transition probabilities $\bar{p}_{01}$ and $\bar{p}_{10}$ defined above model the diffusion of information about the target state s over the social network. We have the following martingale representation theorem for the evolution of Markov process ρ.

Let $\mathcal{F}_k$ denote the sigma algebra generated by $\{\rho_0,\ldots,\rho_k,s_0,\ldots s_k\}$.

Theorem 2.1: For $d = 1,2,\ldots,\bar{D}$, the infected distributions evolve as

$$\begin{aligned}\rho_{k+1}(d) = \rho_k(d) + \frac{1}{N}\big[&\bar{p}_{01}(d,\theta(\rho_k),s_k) \\ &- \bar{p}_{10}(d,\theta(\rho_k),s_k) + w_{k+1}\big] \end{aligned} \quad (11)$$

where w is a martingale increment process, that is $\mathbf{E}\{w_{k+1}|\mathcal{F}_k\} = 0$. Recall s is the finite state Markov chain that models the target process. ■

The above theorem is a well-known martingale representation of a Markov chain [51]—it says that a discrete time Markov process can be obtained by discrete time filtering of a martingale increment process. The theorem implies that the infected distribution dynamics resemble what is commonly called a stochastic approximation (adaptive filtering) algorithm in statistical signal processing: the new estimate is the old estimate plus a noisy update (the "noise" being a martingale increment) that is weighed by a small step size $1/N$ when N is large. Subsequently, we will exploit the structure in Theorem 2.1 to devise a mean field dynamics model which has a state of dimension $\bar{D}$. This is to be compared with the intractable state dimension $\prod_{d=1}^{\bar{D}} N(d)$ of the Markov chain ρ.

C. Mean Field Dynamics of Information Diffusion

The mean field dynamics state that as the number of agents N grows to infinity, the dynamics of the infected distribution ρ, described by (11), in the social network evolves according to the following deterministic difference equation that is modulated by a Markov chain that depends on the target state evolution s:

For $d = 1,2,\ldots,\bar{D}$,

$$\begin{aligned}\bar{\rho}_{k+1}(d) = \bar{\rho}_k(d) + \frac{1}{N}\big[&\bar{p}_{01}(d,\theta(\bar{\rho}_k),s_k) \\ &- \bar{p}_{10}(d,\theta(\bar{\rho}_k),s_k)\big] \end{aligned} \quad (12)$$

$$\bar{p}_{01}(d,\theta,s_k) = (1-\bar{\rho}_k(d)) \sum_{a=0}^{d} p_{01}(d,a,s_k) \binom{d}{a} \theta^a (1-\theta)^{d-a}$$

$$\bar{p}_{10}(d,\theta,s_k) = \bar{\rho}_k(d) \sum_{a=0}^{d} p_{10}(d,a,s_k) \binom{d}{a} \theta^a (1-\theta)^{d-a}$$

$$\theta(\bar{\rho}_k) = \frac{\sum_{d=1}^{\bar{D}} d\,P(d)\,\bar{\rho}_k(d)}{\sum_d^{\bar{D}} d\,P(d)} \quad (13)$$

[2]The transition probabilities are scaled by the degree distribution $P(d)$ for notational convenience. Indeed, since $N(d) = NP(d)$, by using these scaled probabilities we can express the dynamics of the process ρ in terms of the same-step size $1/N$ as described in Theorem 2.1. Throughout this chapter, we assume that the degree distribution $P(d)$, $d \in \{1,2,\ldots,\bar{D}\}$, is uniformly bounded away from zero. That is, $\min_d P(d) > \epsilon$ for some positive constant ϵ.

That the above mean field dynamics follow from (11) is intuitive. Such averaging results are well known in the adaptive filtering community where they are deployed to analyze the convergence of adaptive filters. The difference here is that the limit mean field dynamics are not deterministic but Markov modulated. Moreover, the mean field dynamics here constitute a model for information diffusion, rather than the asymptotic behavior of an adaptive filtering algorithm. As mentioned earlier, from an engineering point of view, the mean field dynamics yield a tractable model for estimation.

We then have the following exponential bound result for the error of the mean field dynamics approximation.

Theorem 2.2: For a discrete time horizon of T points, the deviation between the mean field dynamics $\bar{\rho}_k$ in (13) and actual infected distribution in ρ_k (11) satisfies

$$\mathbb{P}\left\{\max_{0\leq k\leq T}\|\rho_k-\bar{\rho}_k\|_\infty\geq\epsilon\right\}\leq C_1\exp(-C_2\epsilon^2N) \tag{14}$$

where C_1 and C_2 are positive constants and $T=O(N)$. ■

The proof of the above theorem follows from [19, Lemma 1]. Actually in [19] the mean field dynamics are presented in continuous time as a system of ordinary differential equations. The exponential bound follows from an application of the Azuma-Hoeffding inequality. The above theorem provides an exponential bound (in terms of the number of agents N) for the probability of deviation of the sample path of the infected distribution from the mean field dynamics for any finite time interval T.

The stochastic approximation and adaptive filtering literature [52], [53] has several averaging analysis methods for recursions of the form (11). The well studied mean square error analysis [52], [53] computes bounds on $\mathbf{E}\|\bar{\rho}_k-\rho_k\|^2$ instead of the maximum deviation in Theorem 2.2. A mean square error analysis of estimating a Markov modulated empirical distribution is given in [54]. Such mean square analysis assume a finite but small step size $1/N$ in (11).

Numerical Example: We simulate the diffusion of information through a network comprising of $N=100$ nodes (with maximum degree $\bar{D}=17$). It is assumed that at time $k=0$, 5% of nodes are infected. The mean field dynamics model is investigated in terms of the infected link probability (9). The infected link probability $\theta(\rho_k)$ is computed using (13).

Assume each agent is a myopic optimizer and, hence, chooses to adopt the technology only if $c^{(m)}\leq A_k^{(m)}$; $r=1$. At time k, the costs $c^{(m)}$, $m=1,2,\ldots,100$, are i.i.d. random variables simulated from uniform distribution $U[0,C(s_k)]$. Therefore, the transition probabilities in (5) are

$$p_{01}(d,A_k^{(m)},s_k)=\mathbb{P}\left(c^{(m)}\leq A_k^{(m)}\right)=\begin{cases}\frac{A_k^{(m)})}{C(s_k)}, & A_k^{(m)}\leq C(s_k),\\ 1, & A_k^{(m)}>C(s_k).\end{cases} \tag{15}$$

The probability that a product fails is $p_F=0.3$, i.e.,

$$p_{10}(d,A_k^{(m)},s_k)=0.3.$$

The infected link probabilities obtained from network simulation (9) and from the discrete-time mean field dynamics model (13) are illustrated in Figure 2. To illustrate that the infected link probability computed from (13) follows the true one (obtained by network simulation), we assume that the value of C jumps from 1 to 10 at time $k=200$, and from 10 to 1 at time $k=500$. As can be seen in Figure 2, the mean field dynamics provide an excellent approximation to the true infected distribution.

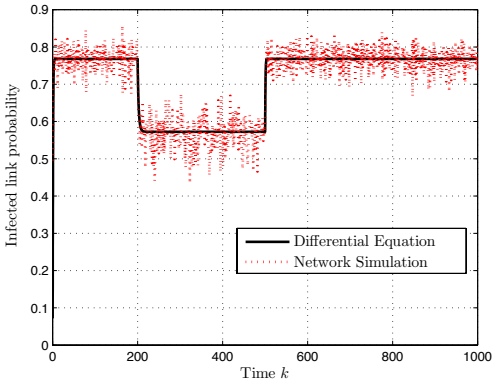

Fig. 2. The infected link probability obtained from network simulation compared to the one obtained from the mean field dynamics model (13). The transition probabilities in (15) depend only on the number of infected neighbors $A_k^{(m)}$ (the parameters are defined in Scenario 1).

D. Example: Social Sensing of Influenza using Twitter

In this section, we utilize datasets from 3 different social networks (namely, (i) Harvard college social network, (ii) influenza datasets from the U.S Centers for Disease Control and Prevention (CDC) and (iii) Twitter, to show how Twitter can be used as a real time social sensor for detecting outbreaks of influenza.

1) Twitter as a social sensor: A key advantage of using social media for rapid sensing of disease outbreaks in health networks is that it is low cost and provides rapid results compared with traditional techniques. For example, CDC must contact thousands of hospitals to query the data which causes a reporting lag of approximately one to two weeks [55]. Using real time microblogging platforms such as Twitter for disease detection has several advantages: the tweets are publicly available, high tweet posting frequency users often provide meta-data (i.e. city, gender, age), and Twitter contains a diverse set of users [55].

Several papers have considered using Twitter data for estimating influenza infection rates. In [56], [57] support vector regression supervised learning algorithms is used to relate the volume of Twitter posts that contain specific words (i.e. *flu, swine, influenza*) to the number of confirmed influenza cases in the U.S. as reported by the CDC. Multiple linear regression [58], [59], and unsupervised Bayesian algorithms [60] have been used to relate the number of tweets of specific words to the influenza rate reported by the CDC. The detection algorithms [56], [58], [60] do not consider the dynamics of the disease propagation and the dynamics of information diffusion in the Twitter network. To reduce the effect of information diffusion in the network, [61] proposes a support vector machine (SVM) classifier to detect: a) if the tweet indicates the users awareness of influenza or indicates the user is infected, and b) if the influenza reference is in reference to another person. The classified tweets are then used to train a multiple linear regression model. To account for the diffusion dynamics of Twitter [62], [63] utilize an Autoregressive with Exogenous input (ARX) model. The exogenous input is the number of unique Twitter users with influenza related tweets, and the output is the number of infected users as reported by the CDC.

If the social network is known then the influenza spread can be formulated in terms of the diffusion model (11). Given the U.S. population of several hundred millions, it is reasonable to adopt the mean field dynamics (13). With the influenza infection rate modelled using (13), the results can be used as an exogenous input to an ARX or Nonlinear ARX (NARX) models to predict the volume of Twitter messages related to influenza as illustrated in Fig.1. In this framework, the Twitter messages are used to validate the underlying propagation model of influenza of use for predicting the infection rate and outbreak detection.

2) Social Network Influenza Dataset: We consider the dataset [64] obtained from a social network of 744 undergraduate students from Harvard College. The health of the 744 students was monitored from September 1, 2009 to December 31, 2009 and was reported by the university Health Services. To construct the social network, students were presented with a background questionnaire. In the questionnaire students are asked: "Please provide the contact information for 2-3 Harvard College students who you know and who you think would like to participate in this study", and "...provide us with the names and contact information of 2-3 of your friends...". This information was used to construct the degree distribution and links of the social network. A movie containing the spread of the influenza in the 744 college students over the 122 day sampling period can be viewed as the Youtube video titled "Social Network Sensors for Early Detection of Contagious Outbreaks" at `http://www.youtube.com/watch?v=0TD06g2m8qM`. Fig.3(a-c) display 3 illustrative snapshots from this video; red nodes denote infected students while yellow nodes depict their neighbors in the social network.

3) Models for Influenza Diffusion: From the data in the youtube video for the Harvard students, we observed the following regarding the transition probabilities $p_{ij}(d, A, s)$ defined in (5). As expected, students with a larger number of infected neighbors A contract influenza sooner. The data shows that the transition probabilities were approximately independent of the degree of the node d. Since the data provided was during an actual influenza outbreak we set the target state of the network (i.e. s) constant. Therefore the transition probabilities depend only on the number of infected neighbours and were estimated as

$$p_{01}(a = 1) = 0.02,\ p_{01}(a = 2) = 0.15.$$

That is, the dataset reveals that the probability of getting infected given $a = 2$ infected neighbors is substantially higher than with $a = 1$ infected neighbor, as expected. The estimated infected link probability θ_k in (9) versus time (days) k is displayed in Fig.3(d). Recall from Sec.II-C that the infected link probability θ_k is related to the mean field dynamics equation (13). This allows the transition probabilities and θ_k to be used to predict the infection rate dynamics.

Other graph-theoretic measures also play a role in the analysis of the diffusion. Students with high k-coreness[3] are expected to contract influenza earlier. Additionally, students that have high betweenness centrality (i.e. number of shortest paths from all students to all

[3] k-coreness is the largest subnetwork comprising nodes of degree at least k

others that pass through that student) contract influenza earlier then students with low betweenness centrality. These observations show that the diffusion of influenza in the network depends strongly on the underlying health network structure. The dynamic model (7) accounts for the effects of the degree of nodes, however to account for the effects from betweenness centrality and k-coreness would require a more sophisticated formulation then that presented in Sec.II-A.

4) Time series model for Influenza Tweets: In Sec.II-D3 we illustrated how the mean field dynamics model (13) can be used to estimate the influenza infection rate with the model parameters estimated from a sampled set of the entire population. To validate the estimated parameters for the entire network requires that the infection rate be related to an observable response, in this case the number of Twitter mentions of a specific keyword. Two time series models are considered for relating the infection rate to the number of Twitter mentions. The models are validated using two real-world datasets of Twitter mentions and number of influenza cases in the U.S..

The number of influenza cases in the U.S. is obtained from the CDC[4] which publishes weekly reports from the U.S. outpatient Influenza-like Illness Surveillance Network (ILInet). The data reported by the CDC is comprised of reports from over 3000 health providers nationwide and was obtained for the dates between September 1, 2012 to October 1, 2013. The associated Twitter data for the 122 day period was obtained using the software PeopleBrowsr[5]. The pre-specified Twitter search terms used are: *flu, swine and influenza*. Since our focus is on monitoring influenza dynamics in the U.S., we excluded all tweets as tagged as originating from outside the U.S. The total number of mentions of a specific keyword on each is obtained using PeopleBrowsr.

We used two time series models for the volume of tweets and compared their performance. The first time series model considered is the ARX model defined by:

$$\tau_k = \sum_{i=1}^{n_a} a_i \tau_{k-i} + \sum_{i=0}^{n_b-1} b_i \rho_{k-\Delta-i} + d + v_k. \qquad (16)$$

In (16), τ_k is the number of influenza related tweets at k, ρ_k is the exogenous input of the infected influenza patients, $n_a, n_b, a_i, b_i, \Delta$ and d are model parameters with v_k an iid noise process. Δ models the delay between patient contraction, and the respective individual tweeting their symptoms. d models the mean number of tweets related to influenza that are not related to an actual infection.

The second time series model we used is the nonlinear autoregressive exogenous (NARX) model given by:

$$\tau_k = F(\tau_{k-1}, \ldots, \tau_{k-n_a}, \rho_{k-\Delta}, \ldots, \rho_{k-\Delta-n_b}) + v_k. \qquad (17)$$

In (17) F denotes a nonlinear function which relates the exogenous input and previous tweets to the current number of tweets. Here we consider F as a support vector machine which can be trained using historical data. Note that if F was independent of previous tweets, previous exogenous inputs, and no delay (i.e. $n_b = 0$ and $\Delta = 0$), then (17) would be identical to the SVM classifier used in [56], [57] to relate the number of tweets to number of infected agents.

The number of reported influenza cases, associated Twitter data, and results of the model training and prediction are displayed in Fig.4 for the ARX (16) and NARX (17) models. As seen from Fig.4(a) ,the dominant word for indicating a possible influenza outbreak is *flu* as compared with *swine* and *influenza*. Notice that there is a lag between the maximum confirmed influenza cases and the # of tweets; however, there is an increase in the number of tweets prior to the peak of infected patients. These dynamics are a result of a combination of infection propagation dynamics and the diffusion of information on Twitter. To account for these dynamics the ARX and NARX models presented in Sec.II-D4 are utilized. The training and prediction accuracy of these models for $n_a = 0, n_b = 2$ (i.e model input parameters $\rho_{k-\Delta}$ and $\rho_{k-\Delta-1}$) are displayed in Fig.4(b). As seen, the NARX (17) model provides a superior estimate as compared with the ARX model (16). Interestingly there is a $\Delta = 18$ day delay between the maximum number of infected patients and the maximum number of Twitter mentions containing the word *flu*. This is contrast to the dynamics observed for the 2009 [56] and 2010-2011 [63] influenza outbreaks which show that the increase in Twitter mentions occurs earlier or at the same time as the number of infected patients increases. This also emphasizes the importance of using the mean field dynamics model for influenza propagation as compared with only using Twitter data for predicting the influenza infection rate. Here we have used the CDC data to estimate the number of infected agents, however the mean field dynamics model (13) could be used to estimate the dynamics of disease propagation and relate this to the observable number of tweets in real-time.

To summarize, the above datasets illustrate how Twitter can be used as a sensor for monitoring the spread of influenza in a heath network. The propagation of

[4]http://gis.cdc.gov/grasp/fluview/fluportaldashboard.html
[5]http://gr.peoplebrowsr.com/

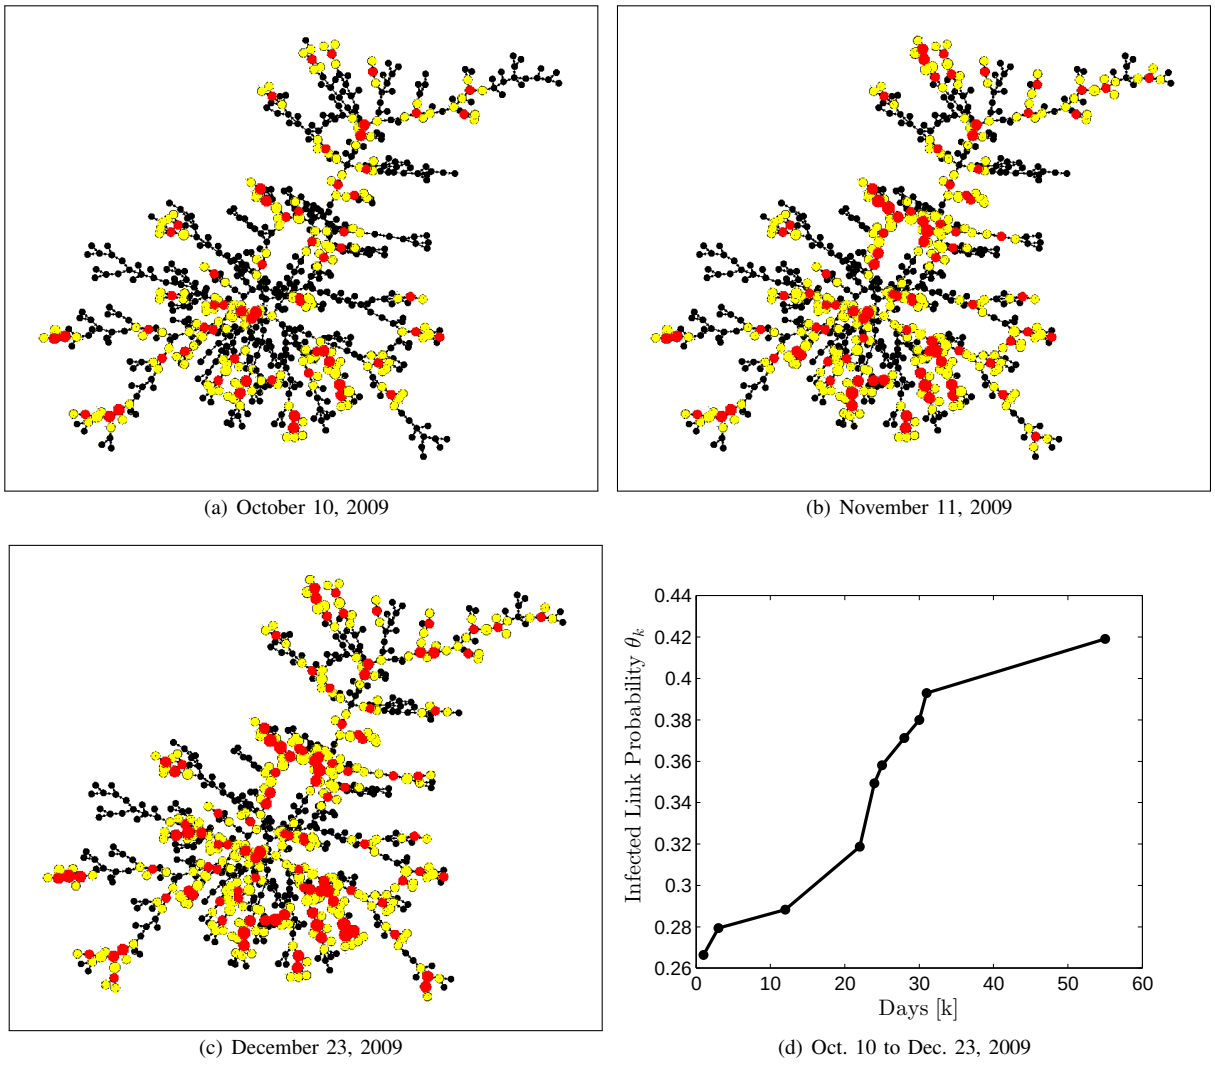

(a) October 10, 2009

(b) November 11, 2009

(c) December 23, 2009

(d) Oct. 10 to Dec. 23, 2009

Fig. 3. Snapshots from youtube video of Harvard undergraduate social network propagation of influenza and the estimated infected link probability θ_k (9) for October 10 to December 23, 2009.

influenza was modeled according to the SIS model and the dynamics of tweets according to an autoregressive model.

E. Sentiment-Based Sensing Mechanism

In the above dataset, samples of influenza affected individuals were obtained from a Harvard college social network. More generally, it is often necessary to sample individuals in a social network to estimate an underlying state of nature such as the sentiment. An important question regarding sensing in a social network is: How can one construct a small but representative sample of a social network with a large number of nodes? In [65] several scale-down and back-in-time sampling procedures are studied. Below we review three sampling schemes. The simplest possible sampling scheme is uniform sampling. We also briefly describe *social sampling* and *respondent-driven sampling* which are recent methods that have become increasingly popular.

1) Uniform Sampling: Consider the following sampling-based measurement strategy. At each period k, $\alpha(d)$ individuals are sampled[6] independently and uniformly from the population $N(d)$ comprising of agents with connectivity degree d. That is, a uniform distributed i.i.d. sequence of nodes, denoted by$\{m_l, l = 1 : \alpha(d)\}$, is generated from the population

[6]For large population sizes N, sampling with and without replacement are equivalent.

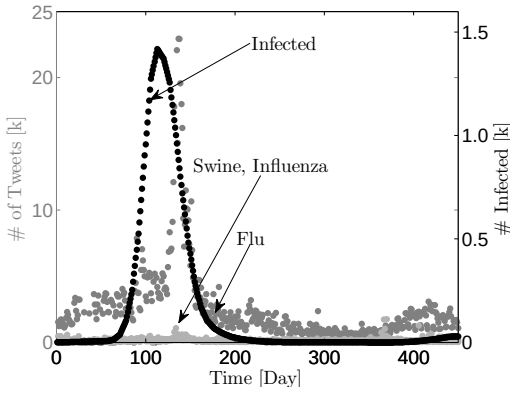

(a) Experimental data of # of tweets and # of Infected with influenza.

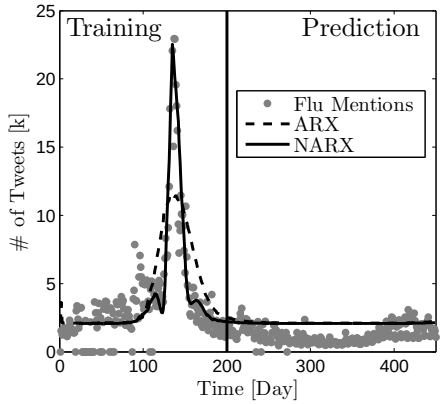

(b) ARX model (16) and NARX model (17) for the # of tweets given the # of infected influenza patients. The ARX and NARX models are trained using the initial 200 days of data. The predictive accuracy of the models is illustrated for the remaining 230 days.

Fig. 4. The experimental data is obtained for September 2012 to October 2013 as described in Sec.II-D2. The ARX (16) and NARX (17) models are utilized to estimate the number of tweets with *flu* given the number of infected influenza patients.

$N(d)$. The messages $y_k^{(m_l)}$ of these $\alpha(d)$ individuals are recorded. From these independent samples, the empirical sentiment distribution $z_k(d)$ of degree d nodes at each time k is obtained as

$$z_k(d,y) = \frac{1}{\alpha(d)} \sum_{l=1}^{\alpha(d)} I\left\{y_k^{(m_l)} = y\right\}, \quad y = 1, \ldots, Y. \tag{18}$$

At each time k, the empirical sentiment distribution z_k can be viewed as noisy observations of the infected distribution ρ_k and target state process s_k.

2) Social Sampling: Social sampling is an extensive area of research; see [66] for recent results. In social sampling, participants in a poll respond with a summary of their friend's responses. This leads to a reduction in the number of samples required. If the average degree of nodes in the network is d, then the savings in the number of samples is by a factor of d, since a randomly chosen node summarizes the results form d of its friends. However, the variance and bias of the estimate depend strongly on the social network structure[7]. In [66], a social sampling method is introduced and analyzed where nodes of degree d are sampled with probability proportional to $1/d$. This is intuitive since weighing neighbors' values by the reciprocal of the degree undoes the bias introduced by large degree nodes. It then illustrates this social sampling method and variants on the LIVEJOURNAL network (livejournal.com) comprising of more than 5 million nodes and 160 million directed edges.

3) MCMC Based Respondent-Driven Sampling (RDS): Respondent-driven sampling (RDS) was introduced by Heckathorn [67], [68], [69] as an approach for sampling from hidden populations in social networks and has gained enormous popularity in recent years. There are more than 120 RDS studies worldwide involving sex workers and injection drug users [70]. As mentioned in [71], the U.S. Centers for Disease Control and Prevention (CDC) recently selected RDS for a 25-city study of injection drug users that is part of the National HIV Behavioral Surveillance System [72].

RDS is a variant of the well known method of snowball sampling where current sample members recruit future sample members. The RDS procedure is as follows: A small number of people in the target population serve as seeds. After participating in the study, the seeds recruit other people they know through the social network in the target population. The sampling continues according to this procedure with current sample members recruiting the next wave of sample members until the desired sampling size is reached. Typically, monetary compensations are provided for participating in the data collection and recruitment.

RDS can be viewed as a form of Markov Chain Monte Carlo (MCMC) sampling (see [71] for an excellent exposition). Let $\{m_l, l = 1 : \alpha(d)\}$ be the realization of an aperiodic irreducible Markov chain with state space $N(d)$ comprising of nodes of degree d. This

[7]In [66], nice intuition is provided in terms of intent polling and expectation polling. In intent polling, individual are sampled and asked who they intend to vote for. In expectation polling, individuals are sampled and asked who they think would win the election. For a given sample size, one would believe that expectation polling is more accurate than intent polling since in expectation polling, an individual would typically consider its own intent together with the intents of its friends.

Markov chain models the individuals of degree d that are snowball sampled, namely, the first individual m_1 is sampled and then recruits the second individual m_2 to be sampled, who then recruits m_3 and so on. Instead of the independent sample estimator (18), an asymptotically unbiased MCMC estimate is then generated as

$$\frac{\sum_{l=1}^{\alpha(d)} \frac{I(y_k^{(m_l)}=y)}{\pi(m_l)}}{\sum_{l=1}^{\alpha(d)} \frac{1}{\pi(m_l)}} \tag{19}$$

where $\pi(m)$, $m \in N(d)$, denotes the stationary distribution of the Markov chain. For example, a reversible Markov chain with prescribed stationary distribution is straightforwardly generated by the Metropolis Hastings algorithm.

In RDS, the transition matrix and, hence, the stationary distribution π in the estimator (19) is specified as follows: Assume that edges between any two nodes m and n have symmetric weights $W(m,n)$ (i.e., $W(m,n) = W(n,m)$, equivalently, the network is undirected). In RDS, node m recruits node n with transition probability $W(m,n)/\sum_n W(m,n)$. Then, it can be easily seen that the stationary distribution is $\pi(m) = \sum_{n\in V} W(m,n)/\sum_{m\in V, n\in V} W(m,n)$. Using this stationary distribution, along with the above transition probabilities for sampling agents in (19), yields the RDS algorithm.

It is well known that a Markov chain over a non-bipartite connected undirected network G is aperiodic. Then, the initial seed for the RDS algorithm can be picked arbitrarily, and the above estimator is an asymptotically unbiased estimator.

Note the difference between RDS and social sampling: RDS uses the network to recruit the next respondent, whereas social sampling seeks to reduce the number of samples by using people's knowledge of their friends' (neighbors') opinions.

Finally, the reader may be familiar with the DARPA network challenge in 2009 where the locations of 10 red balloons in the continental US were to be determined using social networking. In this case, the winning MIT Red Balloon Challenge Team used a recruitment based sampling method. The strategy can also be viewed as a variant of the Query Incentive Network model of [73].

F. Summary and Extensions

This section has discussed the diffusion of information in social networks. Mean field dynamics were used to approximate the asymptotic infected degree distribution. An illustrative example of the spread of influenza was provided. Finally, methods for sampling the population in a social networks were reviews. Below we discuss some related concepts and extensions.

Bayesian Filtering Problem: Given the sentiment observations described above, how can the infected degree distribution ρ_k and target state s_k be estimated at each time instant? The partially observed state space model with dynamics (13) and discrete time observations from sampling the network can be use to obtain Bayesian filtering estimates of the underlying state of nature. Computing the conditional mean estimate s_k, ρ_k given the sentiment observation sequence is a Bayesian filtering problem. In fact, filtering of such jump Markov linear systems have been studied extensively in the signal processing literature [74], [75] and can be solved via the use of sequential Markov chain Monte-Carlo methods. For example, [6] reports on how a particle filter is used to localize earthquake events using Twitter as a social sensor.

Reactive Information Diffusion: A key difference between social sensors and conventional sensors in statistical signal processing is that social sensors are reactive: A social sensor uses additional information gained to modify its behavior. Consider the case where the sentiment-based observation process is made available in a public blog. Then, these observations will affect the transition dynamics of the agents and, therefore, the mean field dynamics.

How Does Connectivity Affect Mean Field Equilibrium?: The papers [46], [17] examine the structure of fixed points of the mean field differential equation (13) when the underlying target process s is not present (equivalently, s is a one state process). They consider the case where the agent transition probabilities are parametrized by $p_{01}(d,a) = \mu F(d,a)$ and $p_{10} = p_F$. Then, defining $\lambda = \mu/p_F$, they study how the following two thresholds behave with the degree distribution and diffusion mechanism:

1) *Critical threshold λ_c:* This is defined as the minimum value of λ for which there exists a fixed point of (13) with positive fraction of infected agents, i.e., $\rho_\infty(d) > 0$ for some d and, for $\lambda \le \lambda_c$, such a fixed point does not exist.
2) *Diffusion threshold λ_d:* Suppose the initial condition ρ_0 for the infected distribution is infinitesimally small. Then, λ_d is the minimum value of λ for which $\rho_\infty(d) > 0$ for some d, and such that, for $\lambda \le \lambda_d$, $\rho_\infty(d) = 0$ for all d.

Determining how these thresholds vary with degree distribution and diffusion mechanism is very useful for understanding the long term behavior of agents in the social network.

III. Bayesian Social Learning Models for Online Reputation Systems

In this section we address the second topic of the paper, namely, Bayesian social learning amongst social sensors. The motivation can be understood in terms of the following social sensing example. Consider the following interactions in a multi-agent social network where agents seek to estimate an underlying state of nature. Each agent visits a restaurant based on reviews on an online reputation website. The agent then obtains a private measurement of the state (e.g., the quality of food in a restaurant) in noise. After that, he reviews the restaurant on the same online reputation website. The information exchange in the social network is modelled by a directed graph. Data incest [76] arises due to loops in the information exchange graph. This is illustrated in the graph of Fig.5. Agents 1 and 2 exchange beliefs (or actions) as depicted in Fig.5. The fact that there are two distinct paths between Agent 1 at time 1 and Agent 1 at time 3 (these paths are denoted in red) implies that the information of Agent 1 at time 1 is double counted leading to a data incest event.

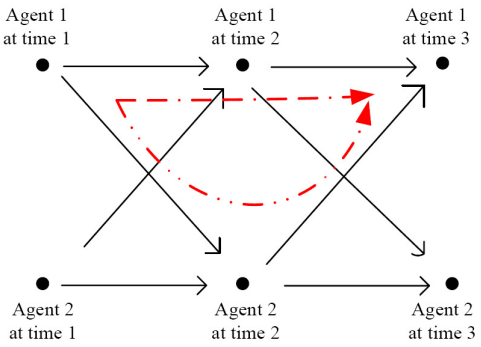

Fig. 5. Example of the information flow (communication graph) in a social network with two agents and over three event epochs. The arrows represent exchange of information.

How can data incest be removed so that agents obtain a fair (unbiased) estimate of the underlying state? The methodology of this section can be interpreted in terms of the recent *Time* article [77] which provides interesting rules for online reputation systems. These include: (i) review the reviewers, and (ii) censor fake (malicious) reviewers. The data incest removal algorithm proposed in this paper can be viewed as "reviewing the reviews" of other agents to see if they are associated with data incest or not.

The rest of this section is organized as follows:

1) Sec.III-A to III-C describe the social learning model that is used to mimic the behavior of agents in online reputation systems. The information exchange between agents in the social network is formulated on a family of time dependent directed acyclic graphs.
2) Sec.III-D presents an incest removal algorithm so that the online reputation system achieves a fair rating. A necessary and sufficient condition is given on the graph structure of information exchange between agents so that a fair rating is achievable.
3) Sec.III-E discusses conditions under which treating individual social sensors as Bayesian optimizers is a useful idealization of their behavior. In particular, it is shown that the ordinal behavior of humans can be mimicked by Bayesian optimizers under reasonable conditions.
4) Sec.III-F presents a dataset obtained from a psychology experiment to illustrate social learning and data incest patterns.

Related work: Collaborative recommendation systems are reviewed and studied in [78], [79]. The books [16], [39] study information cascades in social learning. In [80], a model of Bayesian social learning is considered in which agents receive private information about the state of nature and observe actions of their neighbors in a tree-based network. Another type of mis-information caused by influential agents (agents who heavily affect actions of other agents in social networks) is investigated in [22]. Mis-information in the context of this paper is motivated by sensor networks where the term "data incest" is used [76]. Data incest also arises in Belief Propagation (BP) algorithms [81], [82] which are used in computer vision and error-correcting coding theory. BP algorithms require passing local messages over the graph (Bayesian network) at each iteration. For graphical models with loops, BP algorithms are only approximate due to the over-counting of local messages [83] which is similar to data incest in social learning. With the algorithms presented in this section, data incest can be mitigated from Bayesian social learning over non-tree graphs that satisfy a topological constraint. The closest work to the current paper is [76]. However, in [76], data incest is considered in a network where agents exchange their private belief states - that is, no social learning is considered. Simpler versions of this information exchange process and estimation were investigated in [84], [85], [86]. We also refer the reader to [40] for a discussion of recommender systems.

A. Classical Social Learning

We briefly review the classical social learning model for the interaction of individuals. Subsequently, we will

deal with more general models over a social network.

Consider a multi-agent system that aims to estimate the state of an underlying finite state random variable $x \in \mathbb{X} = \{1, 2, \ldots, X\}$ with known prior distribution π_0. Each agent acts once in a predetermined sequential order indexed by $k = 1, 2, \ldots$ Assume at the beginning of iteration k, all agents have access to the public belief π_{k-1} defined in Step (iv) below. The social learning protocol proceeds as follows [21], [16]:

(i) *Private Observation*: At time k, agent k records a private observation $y_k \in \mathbb{Y}$ from the observation distribution $B_{iy} = P(y|x = i)$, $i \in \mathbb{X}$. Throughout this paper we assume that $\mathbb{Y} = \{1, 2, \ldots, Y\}$ is finite.

(ii) *Private Belief*: Using the public belief π_{k-1} available at time $k-1$ (Step (iv) below), agent k updates its private posterior belief $\eta_k(i) = P(x_k = i|a_1, \ldots, a_{k-1}, y_k)$ using Bayes formula:

$$\eta_k = \frac{B_{y_k}\pi}{\mathbf{1}_X' B_y \pi},\ B_{y_k} = \text{diag}(P(y_k|x = i), i \in \mathbb{X}). \quad (20)$$

Here $\mathbf{1}_X$ denotes the X-dimensional vector of ones, η_k is an X-dimensional probability mass function (pmf).

(iii) *Myopic Action*: Agent k takes action $a_k \in \mathcal{A} = \{1, 2, \ldots, A\}$ to minimize its expected cost

$$a_k = \arg\min_{a\in\mathcal{A}} \mathbf{E}\{c(x, a)|a_1, \ldots, a_{k-1}, y_k\} = \arg\min_{a\in\mathcal{A}} \{c_a'\eta_k\}. \quad (21)$$

Here $c_a = (c(i, a), i \in \mathbb{X})$ denotes an X dimensional cost vector, and $c(i, a)$ denotes the cost incurred when the underlying state is i and the agent chooses action a. Agent k then broadcasts its action a_k.

(iv) *Social Learning Filter*: Given the action a_k of agent k, and the public belief π_{k-1}, each subsequent agent $k' > k$ performs social learning to update the public belief π_k according to the "social learning filter":

$$\pi_k = T(\pi_{k-1}, a_k),\ \text{where } T(\pi, a) = \frac{R_a^\pi \pi}{\sigma(\pi, a)}, \quad (22)$$

where $\sigma(\pi, a) = \mathbf{1}_X' R_a^\pi P'\pi$ is the normalization factor of the Bayesian update. In (22), the public belief $\pi_k(i) = P(x_k = i|a_1, \ldots a_k)$ and $R_a^\pi = \text{diag}(P(a|x = i, \pi), i \in \mathbb{X})$ has elements

$$P(a_k = a|x_k = i, \pi_{k-1} = \pi) = \sum_{y\in\mathbb{Y}} P(a|y, \pi)P(y|x_k = i)$$

$$P(a_k = a|y, \pi) = \begin{cases} 1 \text{ if } c_a' B_y P'\pi \le c_{\tilde{a}}' B_y P'\pi,\ \tilde{a} \in \mathcal{A} \\ 0 \text{ otherwise.} \end{cases}$$

The following result which is well known in the economics literature [21], [16] states that if agents follow the above social learning protocol, then after some finite time $\bar{k}$, an *information cascade* occurs.[8]

Theorem 3.1 ([21]): The social learning protocol leads to an *information cascade* in finite time with probability 1. That is, after some finite time $\bar{k}$ social learning ceases and the public belief $\pi_{k+1} = \pi_k$, $k \ge \bar{k}$, and all agents choose the same action $a_{k+1} = a_k$, $k \ge \bar{k}$. ■

Instead of reproducing the proof, let us give some insight as to why Theorem 3.1 holds. It can be shown using martingale methods that at some finite time $k = k^*$, the agent's probability $P(a_k|y_k, \pi_{k-1})$ becomes independent of the private observation y_k. Then clearly, $P(a_k = a|x_k = i, \pi_{k-1}) = P(a_k = a|\pi)$. Substituting this into the social learning filter (22), we see that $\pi_k = \pi_{k-1}$. Thus after some finite time k^*, the social learning filter hits a fixed point and social learning stops. As a result, all subsequent agents $k > k^*$ completely disregard their private observations and take the same action a_{k^*}, thereby forming an information cascade (and therefore a herd).

B. Data Incest in Online Reputation Systems

In comparison to the previous subsection, we now consider social learning on a family of time dependent directed acyclic graphs - in such cases, apart from herding, the phenomenon of data incest arises.

Consider an online reputation system comprised of agents $\{1, 2, \ldots, S\}$ that aim to estimate an underlying state of nature (a random variable). Let $x \in \mathbb{X} = \{1, 2, \ldots, X\}$ represent the state of nature (such as the quality of a restaurant/hotel) with known prior distribution π_0. Let $k = 1, 2, 3, \ldots$ depict epochs at which events occur. These events involve taking observations, evaluating beliefs and choosing actions as described below. The index k marks the historical order of events. For simplicity, we refer to k as "time".

It is convenient also to reduce the coordinates of time k and agent s to a single integer index n:

$$n \triangleq s + S(k-1), \quad s \in \{1, \ldots, S\},\ k = 1, 2, 3, \ldots \quad (23)$$

We refer to n as a "node" of a time dependent information flow graph G_n which we now define. Let

$$G_n = (V_n, E_n), \quad n = 1, 2, \ldots \quad (24)$$

[8] A *herd of agents* takes place at time $\bar{k}$, if the actions of all agents after time $\bar{k}$ are identical, i.e., $a_k = a_{\bar{k}}$ for all time $k > \bar{k}$. An information cascade implies that a herd of agents occur. [4] quotes the following anecdote of user influence and herding in a social network: "... when a popular blogger left his blogging site for a two-week vacation, the site's visitor tally fell, and content produced by three invited substitute bloggers could not stem the decline."

denote a sequence of time-dependent *directed acyclic graphs* (*DAGs*) [9] of information flow in the social network until and including time k where $n = s+S(k-1)$. Each vertex in V_n represents an agent s' in the social network at time k' and each edge (n', n'') in $E_n \subseteq V_n \times V_n$ shows that the information (action) of node n' (agent s' at time k') reaches node n'' (agent s'' at time k''). It is clear that G_n is a sub-graph of G_{n+1}.

The Adjacency Matrix A_n of G_n is an $n \times n$ matrix with elements $A_n(i,j)$ given by

$$A_n(i,j) = \begin{cases} 1 & \text{if } (v_j, v_i) \in E \\ 0 & \text{otherwise} \end{cases}, \; A_n(i,i) = 0. \tag{25}$$

The transitive closure matrix T_n is the $n \times n$ matrix

$$T_n = \text{sgn}((\mathbf{I}_n - A_n)^{-1}) \tag{26}$$

where for matrix M, the matrix $\text{sgn}(M)$ has elements

$$\text{sgn}(M)(i,j) = \begin{cases} 0 & \text{if } M(i,j) = 0 \\ 1 & \text{if } M(i,j) \neq 0. \end{cases}$$

Note that $A_n(i,j) = 1$ if there is a single hop path between nodes i and j, In comparison, $T_n(i,j) = 1$ if there exists a path (possible multi-hop) between i and j.

The information reaching node n depends on the information flow graph G_n. The following two sets will be used to specify the incest removal algorithms below:

$$\mathcal{H}_n = \{m : A_n(m,n) = 1\} \tag{27}$$

$$\mathcal{F}_n = \{m : T_n(m,n) = 1\}. \tag{28}$$

Thus $\mathcal{H}_n$ denotes the set of previous nodes m that communicate with node n in a single-hop. In comparison, $\mathcal{F}_n$ denotes the set of previous nodes m whose information eventually arrives at node n. Thus $\mathcal{F}_n$ contains all possible multi-hop connections by which information from a node m eventually reaches node n.

Example: Consider $S = 2$ two agents with information flow graph for three time points $k = 1, 2, 3$ depicted in Fig.6 characterized by the family of DAGs $\{G_1, \ldots, G_7\}$. The adjacency matrices $A_1, \ldots, A_7$ are constructed as follows: A_n is the upper left $n \times n$ submatrix of A_{n+1} and

$$A_7 = \begin{bmatrix} 0 & 0 & 1 & 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

[9] A DAG is a directed graph with no directed cycles.

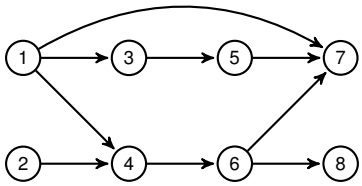

Fig. 6. Example of information flow network with $S = 2$ two agents, namely $s \in \{1, 2\}$ and time points $k = 1, 2, 3, 4$. Circles represent the nodes indexed by $n = s + S(k-1)$ in the social network and each edge depicts a communication link between two nodes.

Let us explain these matrices. Since nodes 1 and 2 do not communicate, clearly A_1 and A_2 are zero matrices. Nodes 1 and 3 communicate, hence A_3 has a single one, etc. Note that if nodes 1,3,4 and 7 are assumed to be the same individual, then at node 7, the individual remembers what happened at node 5 and node 1, but not node 3. This models the case where the individual has selective memory and remembers certain highlights. From (27) and (28),

$$\mathcal{H}_7 = \{1, 5, 6\}, \quad \mathcal{F}_7 = \{1, 2, 3, 4, 5, 6\}$$

where $\mathcal{H}_7$ denotes all one hop links to node 7 while $\mathcal{F}_7$ denotes all multihop links to node 7.

C. Data Incest Model and Social Influence Constraint

Each node n receives recommendations from its immediate friends (one hop neighbors) according to the information flow graph defined above. That is, it receives actions $\{a_m, m \in \mathcal{H}_n\}$ from nodes $m \in \mathcal{H}_n$ and then seeks to compute the associated public beliefs $\pi_m, m \in \mathcal{H}_n$. If node n naively (incorrectly) assumes that the public beliefs $\pi_m, m \in \mathcal{H}_n$ are independent, then it would fuse these as

$$\pi_n - = \frac{\prod_{m \in \mathcal{H}_n} \pi_m}{\mathbf{1}_X' \prod_{m \in \mathcal{H}_n} \pi_m}. \text{ (WRONG!)} \tag{29}$$

This naive data fusion would result in data incest.

1) Aim: The aim is to provide each node n the true posterior distribution

$$\pi^0_{n-}(i) = P(x = i | \{a_m, m \in \mathcal{F}_n\}) \tag{30}$$

subject to the following *social influence constraint*: There exists a fusion algorithm $\mathcal{A}$ such that

$$\pi^0_{n-} = \mathcal{A}(\pi_m, m \in \mathcal{H}_n). \tag{31}$$

2) Discussion. Fair Rating and Social Influence: We briefly pause to discuss (30) and (31).
(i) We call π^0_{n-} in (30) the *true* or *fair online rating* available to node n since $\mathcal{F}_n$ defined in (28) denotes all information (multi-hop links) available to node n. By definition π^0_{n-} is incest free and is the desired conditional probability that agent n needs.

Indeed, if node n combines π^0_{n-} together with its own private observation via social learning, then clearly

$$\eta_n(i) = P(x = i|\{a_m, m \in \mathcal{F}_n\}, y_n), \quad i \in \mathbb{X},$$
$$\pi_n(i) = P(x = i|\{a_m, m \in \mathcal{F}_n\}, a_n), \quad i \in \mathbb{X},$$

are, respectively, the correct (incest free) private belief for node n and the correct after-action public belief. If agent n does not use π^0_{n-}, then incest can propagate; for example if agent n naively uses (29).

Why should an individual n agree to use π^0_{n-} to combine with its private message? It is here that the social influence constraint (31) is important. $\mathcal{H}_n$ can be viewed as the "*social message*", i.e., personal friends of node n since they directly communicate to node n, while the associated beliefs can be viewed as the "*informational message*". As described in the remarkable recent paper [87], the social message from personal friends exerts a large social influence[10] – it provides significant incentive (peer pressure) for individual n to comply with the protocol of combining its estimate with π^0_{n-} and thereby prevent incest. [87] shows that receiving messages from known friends has significantly more influence on an individual than the information in the messages. This study includes a comparison of information messages and social messages on Facebook and their direct effect on voting behavior. To quote [87], "The effect of social transmission on real-world voting was greater than the direct effect of the messages themselves..." In Sec.III-F, we provide results of an experiment on human subjects that also illustrates social influence in social learning. [88] is an influential paper in the area of social influence.

[10]In a study conducted by social networking site *myYearbook*, 81% of respondents said they had received advice from friends and followers relating to a product purchase through a social site; 74 percent of those who received such advice found it to be influential in their decision. (*Click Z*, January 2010).

D. Incest Removal in Online Reputation System

It is convenient to work with the logarithm of the un-normalized belief[11]; accordingly define

$$l_n(i) \propto \log \pi_n(i), \quad l_{n-}(i) \propto \log \pi_{n-}(i), \quad i \in \mathbb{X}.$$

The following theorem shows that the logarithm of the fair rating π^0_{n-} defined in (30) can be obtained as a weighted linear combination of the logarithms of previous public beliefs.

Theorem 3.2 (Fair Rating Algorithm): Suppose the network administrator runs the following algorithm for an online reputation system:

$$l_{n-}(i) = w_n' \, l_{1:n-1}(i)$$
$$\text{where} \;\; w_n = T_{n-1}^{-1} t_n. \tag{32}$$

Then $l_{n-}(i) \propto \log \pi^0_{n-}(i)$. That is, the fair rating $\log \pi^0_{n-}(i)$ defined in (30) is obtained. In (32), w_n is an $n-1$ dimensional weight vector. Recall that t_n denotes the first $n-1$ elements of the nth column of the transitive closure matrix T_n. ■

Theorem 3.2 says that the fair rating π^0_{n-} can be expressed as a linear function of the action log-likelihoods in terms of the transitive closure matrix T_n of graph G_n. This is intuitive since π^0_{n-} can be viewed as the sum of information collected by the nodes such that there are paths between all these nodes and n.

Theorem 3.3 (Achievability of Fair Rating): Consider the fair rating algorithm specified by (32). With available information $(\pi_m, m \in \mathcal{H}_n)$ to achieve the estimates l_{n-} of algorithm (32), a necessary and sufficient condition on the information flow graph G_n is

$$A_n(j, n) = 0 \implies w_n(j) = 0. \tag{33}$$

(Recall w_n is specified in (32).) ■

Note that the constraint (33) is purely in terms of the adjacency matrix A_n, since the transitive closure matrix (26) is a function of the adjacency matrix.

E. Ordinal Decisions and Bayesian Social Sensors

The social learning protocol assumes that each agent is a Bayesian utility optimizer. The following discussion puts together ideas from the economics literature to show that under reasonable conditions, such a Bayesian model

[11]The un-normalized belief proportional to $\pi_n(i)$ is the numerator of the social learning filter (22). The corresponding un-normalized fair rating corresponding to $\pi^0_{n-}(i)$ is the joint distribution $P(x = i, \{a_m, m \in \mathcal{F}_n\})$. By taking the logarithm of the un-normalized belief, Bayes formula merely becomes the sum of the log likelihood and log prior. This allows us to devise a data incest removal algorithm based on linear combinations of the log beliefs.

is a useful idealization of agents' behaviors. This means that the Bayesian social learning follows simple intuitive rules and is therefore, a useful idealization. (In Sec.IV, we discuss the theory of revealed preferences which yields a non-parametric test on data to determine if an agent is a utility maximizer.)

Humans typically make *monotone* decisions - the more favorable the private observation, the higher the recommendation. Humans make *ordinal* decisions[12] since humans tend to think in symbolic ordinal terms. Under what conditions is the recommendation a_n made by node n *monotone increasing* in its observation y_n and *ordinal*? Recall from the social learning protocol (21) that the actions of agents are

$$a_k = \arg\min_{a\in\mathcal{A}}\{c_a' B_{y_k}\pi_k\}.$$

So an equivalent question is: Under what conditions is the argmin increasing in observation y_n? Note that an increasing argmin is an *ordinal* property - that is, $\operatorname{argmin}_a c_a' B_{y_n}\pi_{n-}^0$ increasing in y implies $\operatorname{argmin}_a \phi(c_a' B_{y_n}\pi_{n-}^0)$ is also increasing in y for any monotone function $\phi(\cdot)$.

The following result gives sufficient conditions for each agent to give a recommendation that is monotone and ordinal in its private observation:

Theorem 3.4: Suppose the observation probabilities and costs satisfy the following conditions:

(A1) B_{iy} are TP2 (totally positive of order 2); that is, $B_{i+1,y}B_{i,y+1} \leq B_{i,y}B_{i+1,y+1}$.

(A2) $c(x,a)$ is submodular. That is, $c(x,a+1) - c(x,a) \leq c(x+1,a+1) - c(x+1,a)$.

Then

1) Under (A1) and (A2), the recommendation $a_n(\pi_{n-}^0, y_n)$ made by agent n is increasing and hence ordinal in observation y_n, for any π_{n-}^0.
2) Under (A2), $a_n(\pi_{n-}^0, y_n)$ is increasing in belief π_{n-}^0 with respect to the monotone likelihood ratio (MLR) stochastic order[13] for any observation y_n.■

The proof is in [89]. We can interpret the above theorem as follows. If agents makes recommendations that are monotone and ordinal in the observations and monotone in the prior, then they mimic the Bayesian social learning model. Even if the agent does not exactly follow a Bayesian social learning model, its monotone ordinal behavior implies that such a Bayesian model is a useful idealization.

Condition (A1) is widely studied in monotone decision making; see the classical book by Karlin [90] and [91]; numerous examples of noise distributions are TP2. Indeed in the highly cited paper [92] in the economics literature, observation $y+1$ is said to be more "favorable news" than observation y if Condition (A1) holds.

Condition (A2) is the well known submodularity condition [93], [94], [95]. (A2) makes sense in a reputation system for the costs to be well posed. Suppose the recommendations in action set $\mathcal{A}$ are arranged in increasing order and also the states in $\mathbb{X}$ for the underlying state are arranged in ascending order. Then (A2) says: if recommendation $a+1$ is more accurate than recommendation a for state x; then recommendation $a+1$ is also more accurate than recommendation a for state $x+1$ (which is a higher quality state than x).

In the experiment results reported in Sec.III-F, we found that (A1) and (A2) of Theorem 3.4 are justified.

F. Psychology Experiment Dataset

To illustrate social learning, data incest and social influence, this section presents an actual psychology experiment that was conducted by our colleagues at the Department of Psychology of University of British Columbia in September and October, 2013. The participants comprised 36 undergraduate students who participated in the experiment for course credit.

1) Experiment Setup: The experimental study involved 1658 individual trials. Each trial comprised two participants who were asked to perform a perceptual task interactively. The perceptual task was as follows: Two arrays of circles denoted left and right, were given to each pair of participants. Each participant was asked to judge which array (left or right) had the larger average diameter. The participants answer (left of right) constituted their action. So the action space is $\mathcal{A} = \{0 \text{ (left)}, 1 \text{ (right)}\}$.

The circles were prepared for each trial as follows: two 4×4 grids of circles were generated by uniformly sampling from the radii: $\{20, 24, 29, 35, 42\}$ (in pixels). The average diameter of each grid was computed, and if the means differed by more than 8% or less than 4%, new grids were made. Thus in each trial, the left array and right array of circles differed in the average diameter by 4-8%

For each trial, one of the two participants was chosen randomly to start the experiment by choosing an action according to his/her observation. Thereafter, each participant was given access to their partner's previous

[12]Humans typically convert numerical attributes to ordinal scales before making decisions. For example, it does not matter if the cost of a meal at a restaurant is \$200 or \$205; an individual would classify this cost as "high". Also credit rating agencies use ordinal symbols such as AAA, AA, A.

[13] Given probability mass functions $\{p_i\}$ and $\{q_i\}$, $i = 1, \ldots, X$ then p MLR dominates q if $\log p_i - \log p_{i+1} \leq \log q_i - \log q_{i+1}$.

response (action) and the participants own previous action prior to making his/her judgement. This mimics the social learning protocol. The participants continued choosing actions according to this procedure until the experiment terminated. The trial terminated when the response of each of the two participants did not change for three successive iterations (the two participants did not necessarily have to agree for the trial to terminate).

In each trial, the actions of participants were recorded along with the time interval taken to choose their action. As an example, Fig. 8 illustrates the sample path of decisions made by the two participants in one of the 1658 trials. In this specific trial, the average diameter of the left array of circles was 32.1875 and the right array was 30.5625 (in pixels); so the ground truth was 0 (left).

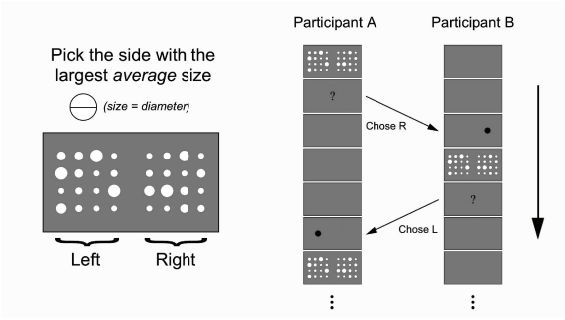

Fig. 7. Two arrays of circles were given to each pair of participants on a screen. Their task is to interactively determine which side (either left or right) had the larger average diameter. The partner's previous decision was displayed on screen prior to the stimulus.

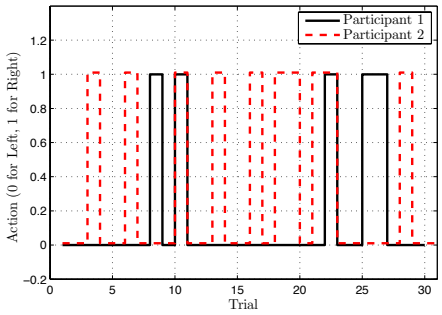

Fig. 8. Example of sample path of actions chosen by two participants in a single trial of the experiment. In this trial, both participants eventually chose the correct answer 0 (left).

2) Experimental Results: The results of our experimental study are as follows:

a) Social learning Model: As mentioned above, the experiment for each pair of participants was continued until both participants' responses stabilized. *In what percentage of these experiments, did an agreement occur between the two participants?* The answer to this question reveals whether "herding" occurred in the experiments and whether the participants performed social learning (influenced by their partners). The experiments show that in 66% of trials (1102 among 1658), participants reached an agreement; that is herding occurred. Further, in 32% of the trials, both participants converged to the correct decision after a few interactions.

To construct a social learning model for the experimental data, we consider the experiments where both participants reached an agreement. Define the social learning success rate as

$$\frac{\text{\# expts where both participants chose correct answer}}{\text{\# expts where both participants reached an agreement}}.$$

In the experimental study, the state space is $\mathbb{X} = \{0, 1\}$ where $x = 0$, when the left array of circles has the larger diameter and $x = 1$, when the right array has the larger diameter. The initial belief for both participants is considered to be $\pi_0 = [0.5, 0.5]$. The observation space is assumed to be $\mathbb{Y} = \{0, 1\}$.

To estimate the social learning model parameters (observation probabilities B_{iy} and costs $c(i, a)$), we determined the parameters that best fit the learning success rate of the experimental data. The best fit parameters obtained were[14]

$$B_{iy} = \begin{bmatrix} 0.61 & 0.39 \\ 0.41 & 0.59 \end{bmatrix}, \quad c(i, a) = \begin{bmatrix} 0 & 2 \\ 2 & 0 \end{bmatrix}.$$

Note that B_{iy} and $c(i, a)$ satisfy both the conditions of the Theorem 3.4, namely TP2 observation probabilities and single-crossing cost. This implies that the subjects of this experiment made monotone and ordinal decisions.

b) Data incest: Here, we study the effect of information patterns in the experimental study that can result in data incest. Since private observations are highly subjective and participants did not document these, we cannot claim with certainty if data incest changed the action of an individual. However, from the experimental data, we can localize specific information patterns that can result in incest. In particular, we focus on the two information flow graphs depicted in Fig.9. In the two graphs of Fig.9, the action of the first participant at time k influenced the action of the second participant at time $k + 1$, and thus, could have been double counted by the first participant at time $k + 2$. We found that in 79% of

[14]Parameter estimation in social learning is a challenging problem not addressed in this paper. Due to the formation of cascades in finite time, construction of an asymptotically consistent estimator is impossible, since actions after the formation of a cascade contain no information.

experiments, one of the information patterns shown in Fig.9 occurred (1303 out of 1658 experiments). Further, in 21% of experiments, the information patterns shown in Fig.9 occurred and at least one participant changed his/her decision, i.e., the judgement of participant at time $k+1$ differed from his/her judgements at time $k+2$ and k. These results show that even for experiments involving two participants, data incest information patterns occur frequently (79%) and causes individuals to modify their actions (21%). It shows that social learning protocols require careful design to handle and mitigate data incest.

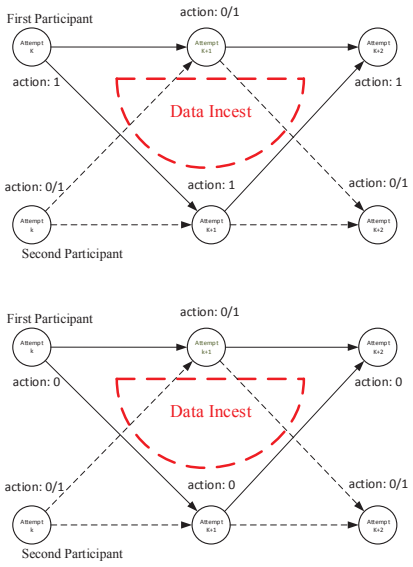

Fig. 9. Two information patterns from our experimental studies which can result in data incest.

G. Summary and Extensions

In this section, we have outlined a controlled sensing problem over a social network in which the administrator controls (removes) data incest and thereby maintains an unbiased (fair) online reputation system. The state of nature could be geographical coordinates of an event (in a target localization problem) or quality of a social unit (in an online reputation system). As discussed above, data incest arises due to the recursive nature of Bayesian estimation and non-determinism in the timing of the sensing by individuals. Details of proofs, extensions and further numerical studies are presented in [76], [96].

We summarize some extensions of the social learning framework that are relevant to interactive sensing.

1) Wisdom of Crowds: Surowiecki's book [97] is an excellent popular piece that explains the wisdom-of-crowds hypothesis. The wisdom-of-crowds hypothesis predicts that the independent judgments of a crowd of individuals (as measured by any form of central tendency) will be relatively accurate, even when most of the individuals in the crowd are ignorant and error prone. The book also studies situations (such as rational bubbles) in which crowds are not wiser than individuals. Collect enough people on a street corner staring at the sky, and everyone who walks past will look up. Such herding behavior is typical in social learning.

2) In which order should agents act?: In the social learning protocol, we assumed that the agents act sequentially in a pre-defined order. However, in many social networking applications, it is important to optimize the order in which agents act. For example, consider an online review site where individual reviewers with different reputations make their reviews publicly available. If a reviewer with high reputation publishes her review first, this review will unduly affect the decision of a reviewer with lower reputation. In other words, if the most senior agent "speaks" first it would unduly affect the decisions of more junior agents. This could lead to an increase in bias of the underlying state estimate.[15] On the other hand, if the most junior agent is polled first, then since its variance is large, several agents would need to be polled in order to reduce the variance. We refer the reader to [99] for an interesting description of who should speak first in a public debate.[16] It turns out that for two agents, the seniority rule is always optimal for any prior – that is, the senior agent speaks first followed by the junior agent; see [99] for the proof. However, for more than two agents, the optimal order depends on the prior, and the observations in general.

IV. Revealed Preferences: Are Social Sensors Utility Maximizers?

We now move on to the third main topic of the paper, namely, the principle of revealed preferences. The main question addressed is: Given a dataset of decisions made by a social sensor, is it possible to determine if the social

[15]To quote a recent paper from the Haas School of Business, U.C. Berkeley [98]: "In 94% of cases, groups (of people) used the first answer provided as their final answer... Groups tended to commit to the first answer provided by any group member." People with dominant personalities tend to speak first and most forcefully "even when they actually lack competence".

[16]As described in [99], seniority is considered in the rules of debate and voting in the U.S. Supreme Court. "In the past, a vote was taken after the newest justice to the Court spoke, with the justices voting in order of ascending seniority largely, it was said, to avoid the pressure from long-term members of the Court on their junior colleagues."

sensor is a utility maximizer? More generally, is a dataset from a multiagent system consistent with play from a Nash equilibrium? If yes, can the behavior of the social sensors be learned using data from the social network?

These questions are fundamentally different to the model-based theme that is widely used in the communications literature in which an objective function (typically convex) is proposed and then algorithms are constructed to compute the minimum. In contrast, the revealed preference approach is data-centric–we wish to determine whether the dataset is obtained from an utility maximizer. In simple terms, revealed preference theory seeks to determine if an agent is an utility maximizer subject to budget constraints based on observing its choices over time. The principle of revealed preferences is widely studied in the micro-economics literature. As mentioned in Sec.I, Varian (chief economist at Google) has written several influential papers in this area. In this section we will use the principle of revealed preferences on datasets to determine how social sensors behave as a function of external influence. The setup is depicted in the schematic diagram Fig.10.

A. Afriat's Theorem for a single agent

The theory of revealed preferences was pioneered by Samuelson [100]. Afriat published a highly influential paper [29] in revealed preferences (see also [101]). Given a time-series of data $\mathcal{D} = \{(p_t, x_t), t \in \{1, 2, \ldots, T\}\}$ where $p_t \in \mathbb{R}^m$ denotes the external influence, x_t denotes the response of agent, and t denotes the time index, is it possible to detect if the agent is a *utility maximizer*? An agent is a *utility maximizer* if for every external influence p_t, the chosen response x_t satisfies

$$x_t(p_t) \in \underset{\{p_t' x \leq I_t\}}{\arg\max}\, u(x) \tag{34}$$

with $u(x)$ a non-satiated utility function. Nonsatiated means that an increase in any element of response x results in the utility function increasing.[17] As shown by Diewert [102], without local nonsatiation the maximization problem (34) may have no solution.

In (34) the budget constraint $p_t' x \leq I_t$ denotes the total amount of resources available to the social sensor for selecting the response x to the external influence p_t. For example, if p_t is the electricity price and x_t the associated electricity consumption, then the budget of the social sensor is the available monetary funds for purchasing electricity. In the real-world social sensor datasets provided in this paper, further insights are provided for the budget constraint.

The celebrated "Afriat's theorem" provides a necessary and sufficient condition for a finite dataset $\mathcal{D}$ to have originated from an utility maximizer. Afriat's theorem has subsequently been expanded and refined, most notably by Diewert [102], Varian [30] and Blundell [103].

Theorem 4.1 (Afriat's Theorem): Given a dataset $\mathcal{D} = \{(p_t, x_t) : t \in \{1, 2, \ldots, T\}\}$, the following statements are equivalent:

1) The agent is a utility maximizer and there exists a nonsatiated and concave utility function that satisfies (34).
2) For scalars u_t and $\lambda_t > 0$ the following set of inequalities has a feasible solution:
$$u_\tau - u_t - \lambda_t p_t'(x_\tau - x_t) \leq 0 \text{ for } t, \tau \in \{1, 2, \ldots, T\}. \tag{35}$$
3) A nonsatiated and concave utility function that satisfies (34) is given by:
$$u(x) = \min_{t \in T} \{u_t + \lambda_t p_t'(x - x_t)\} \tag{36}$$
4) The dataset $\mathcal{D}$ satisfies the Generalized Axiom of Revealed Preference (GARP), namely for any $k \leq T$, $p_t' x_t \geq p_t' x_{t+1} \quad \forall t \leq k-1 \implies p_k' x_k \leq p_k' x_1$.

■

As pointed out in Varian's influential paper [30], a remarkable feature of Afriat's theorem is that if the dataset can be rationalized by a non-trivial utility function, then it can be rationalized by a continuous, concave, monotonic utility function. "Put another way, violations of continuity, concavity, or monotonicity cannot be detected with only a finite number of demand observations".

Verifying GARP (statement 4 of Theorem 4.1) on a dataset $\mathcal{D}$ comprising T points can be done using Warshall's algorithm with $O(T^3)$ [30], [104] computations. Alternatively, determining if Afriat's inequalities (35) are feasible can be done via a LP feasibility test (using for example interior point methods [105]). Note that the utility (36) is not unique and is ordinal by construction. Ordinal means that any monotone increasing transformation of the utility function will also satisfy Afriat's theorem. Therefore the utility mimics the ordinal behavior of humans, see also Sec.III-E. Geometrically the estimated utility (36) is the lower envelop of a finite number of hyperplanes that is consistent with the dataset $\mathcal{D}$.

Note that GARP is equivalent to the notion of "cyclical consistency" [106] – they state that the responses are consistent with utility maximization if no negative cycles are present. As an example, consider a dataset $\mathcal{D}$ with

[17] The non-satiated assumption rules out trivial cases such as a constant utility function which can be optimized by any response.

$T = 2$ observations resulting from a utility maximization agent. Then GARP states that $p_1'x_1 \geq p_1'x_2 \implies p_2'x_2 \leq p_2'x_1$. From (34), the underlying utility function must satisfy $u(x_1) \geq u(x_2) \implies u(x_2) \leq u(x_1)$ where the equality results if $x_1 = x_2$.

Another remarkable feature of Afriat's Theorem is that no parametric assumptions of the utility function of the agent are necessary. To gain insight into the construction of the inequalities (35), let us assume the utility function $u(x)$ in (34) is increasing for positive x, concave, and differentiable. If x_t solves the maximization problem (34), then from the Karush-Kuhn-Tucker (KKT) conditions there must exist Lagrange multipliers λ_t such that

$$\nabla u(x_t) = \lambda_t \nabla(p_t'x_t - I_t) = \lambda_t p_t$$

is satisfied for all $t \in \{1, 2, \ldots, T\}$. Note that since $u(x)$ is increasing, $\nabla u(x_t) = \lambda_t p_t > 0$, and since p_t is strictly positive, $\lambda_t > 0$. Given that $u(x)$ is a concave differentiable function, it follows that

$$u(x) \leq u(x_t) + \nabla u(x_t)'(x - x_t) \quad \forall t \in \{1, 2, \ldots, T\}.$$

Denoting $u_t = u(x_t)$ and $u_\tau = u(x_\tau)$, and using the KKT conditions and concave differential property, the inequalities (35) result. To prove that if the solution of (35) is feasible then GARP is satisfied can be performed using the duality theorem of linear programming as illustrated in [104].

B. Revealed Preferences for Multi-agent Social Sensors

We now consider a multi-agent version of Afriat's theorem for deciding if a dataset is generated by playing from the equilibrium of a potential game[18] An example is the control of power consumption in the electrical grid. Consider a corporate network of financial management operators that select the electricity prices in a set of zones in the power grid. By selecting the prices of electricity the operators are expected to be able to control the power consumption in each zone. The operators wish to supply their consumers with sufficient power however given the finite amount of resources the operators in the corporate network must interact. This behavior can be modelled as a game. Recent analysis of energy use scheduling and demand side management schemes in the energy market have been performed using potential games [108], [109], [110]. Another example of potential games are *congestion games* [111], [112], [113], [114] in which the utility of each player depends on the amount of resource it and other players use.

[18]As in [107], we consider potential games since they are sufficiently specialized so that there exist datasets that fail Afriat's test.

Consider the social network presented in Fig.10, given a time-series of data from N agents $\mathcal{D} = \{(p_t, x_t^1, \ldots, x_t^n) : t \in \{1, 2, \ldots, T\}\}$ with $p_t \in \mathbb{R}^m$ the external influence, x_t^i the response of agent i, and t the time index, is it possible to detect if the dataset originated from agents that play a potential game?

The characterization of how agents behave as a function of external influence, for example price of using a resource, and the responses of other agents in a social network, is key for analysis. Consider the social network illustrated in Fig.10. There are a total of n interacting agents in the network and each can produce a response x_t^i in response to the other agents and an external influence p_t. Without any *a priori* assumptions about the agents, how can the behaviour of the agents in the social network be learned? In the engineering literature the behaviour of agents is typically defined *a priori* using a *utility function*, however our focus here is on learning the behaviour of agents. The *utility function* captures the satisfaction or payoff an agent receives from a set of possible responses, denoted by X. Formally, a utility function $u : X \to \mathbb{R}$ represents a preference relation between responses x_1 and x_2 if and only if for every $x_1, x_2 \in X$, $u(x_1) \leq u(x_2)$ implies x_2 is preferred to x_1. Given a time-series of data $\mathcal{D} = \{(p_t, x_t^1, \ldots, x_t^n) : t \in \{1, 2, \ldots, T\}\}$ with $p_t \in \mathbb{R}^m$ the external influence, x_t^i the response of agent i, and t the time index, is it possible to detect if the series originated from an agent that is a *utility maximizer*?

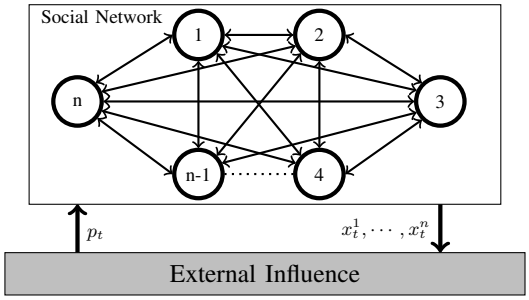

Fig. 10. Schematic of a social network containing n agents where $p_t \in \mathbb{R}^m$ denotes the external influence, and $x_t^i \in \mathbb{R}^m$ the response of agent i in response to the external influence and other agents at time t. Note that dotted line denotes consumers $4, \ldots, n-1$. The aim is to determine if the dataset $\mathcal{D}$ defined in (37), is consistent with play from a Nash equilibrium.

In a network of social sensors (Fig.10), the responses of agents may be dependent on both the external influence p_t and the responses of the other agents in the network, denoted by x_t^{-i}. The utility function of the agent must now include the responses of other agents–formally if there are n agents, each has a utility

function $u^i(x^i, x_t^{-i})$ with x^i denoting the response of agent i, x_t^{-i} the responses of the other $n-1$ agents, and $u^i(\cdot)$ the utility of agent i. Given a dataset $\mathcal{D}$, is it possible to detect if the data is consistent with agents playing a game and maximizing their individual utilities? Deb, following Varian's and Afriat's work, shows that refutable restrictions exist for the dataset $\mathcal{D}$, given by (37), to satisfy Nash equilibrium (38) [107], [29], [115]. These refutable restrictions are however, satisfied by most $\mathcal{D}$ [107]. The detection of agents engaged in a concave potential game, and generating responses that satisfy Nash equilibrium, provide stronger restrictions on the dataset $\mathcal{D}$ [107], [116]. We denote this behaviour as *Nash rationality*, defined as follows:

Definition 4.1 ([116], [117], [118]): Given a dataset

$$\mathcal{D} = \{(p_t, x_t^1, x_t^2, \ldots, x_t^n) : t \in \{1, 2, \ldots, T\}\}, \quad (37)$$

$\mathcal{D}$ is consistent with *Nash equilibrium* play if there exist utility functions $u^i(x^i, x^{-i})$ such that

$$x_t^i = x_t^{i*}(p_t) \in \underset{\{p_t' x^i \leq I_t^i\}}{\arg\max}\, u^i(x^i, x^{-i}). \quad (38)$$

In (38), $u^i(x, x^{-i})$ is a nonsatiated utility function in x, $x^{-i} = \{x^j\}_{j \neq i}$ for $i, j \in \{1, 2, \ldots, n\}$, and the elements of p_t are strictly positive. Nonsatiated means that for any $\epsilon > 0$, there exists a x^i with $\|x^i - x_t^i\|_2 < \epsilon$ such that $u^i(x^i, x^{-i}) > u^i(x_t^i, x_t^{-i})$. If for all $x^i, x^j \in X^i$, there exists a concave potential function V that satisfies

$$u^i(x^i, x^{-i}) - u^i(x^j, x^{-i}) > 0 \quad \text{iff } V(x^i, x^{-i}) - V(x^j, x^{-i}) > 0 \quad (39)$$

for all the utility functions $u^i(\cdot)$ with $i \in \{1, 2, \ldots, n\}$, then the dataset $\mathcal{D}$ satisfies *Nash rationality*. ■

Just as with the utility maximization budget constraint in (34), the budget constraint $p_t' x^i \leq I_t^i$ in (38) models the total amount of resources available to the social sensor for selecting the response x_t^i to the external influence p_t.

The detection test for Nash rationality (Definition 4.1) has been used in [119] to detect if oil producing countries are collusive, and in [107] for the analysis of household consumption behaviour.

In the following sections, decision tests for utility maximization, and a non-parametric learning algorithm for predicting agent responses are presented. Three real world datasets are analyzed using the non-parametric decision tests and learning algorithm. The datasets are comprised of bidders auctioning behaviour, electrical consumption in the power grid, and on the tweeting dynamics of agents in the social network Twitter illustrated in Fig.10.

C. Decision Test for Nash Rationality

This section presents a non-parametric test for Nash rationality given the dataset $\mathcal{D}$ defined in (37). If the dataset $\mathcal{D}$ passes the test, then it is consistent with play according to a Nash equilibrium of a concave potential game. In Sec.IV-D, a learning algorithm is provided that can be used to predict the response of agents in the social network provided in Fig.10.

The following theorem provides necessary and sufficient conditions for a dataset $\mathcal{D}$ to be consistent with Nash rationality (Definition 4.1). The proof is analogous to Afriat's Theorem when the concave potential function of the game is differentiable [107], [116], [120].

Theorem 4.2 (Multiagent Afriat's Theorem): Given a dataset $\mathcal{D}$ (37), the following statements are equivalent:

1) $\mathcal{D}$ is consistent with Nash rationality (Definition 4.1) for an n-player concave potential game.
2) Given scalars v_t and $\lambda_t^i > 0$ the following set of inequalities have a feasible solution for $t, \tau \in \{1, \ldots, T\}$,
$$v_\tau - v_t - \sum_{i=1}^{n} \lambda_t^i p_t'(x_\tau^i - x_t^i) \leq 0. \quad (40)$$
3) A concave potential function that satisfies (38) is given by:
$$\hat{V}(x^1, x^2, \ldots, x^n) = \min_{t \in T} \{v_t + \sum_{i=1}^{n} \lambda_t^i p_t'(x^i - x_t^i)\}. \quad (41)$$
4) The dataset $\mathcal{D}$ satisfies the Potential Generalized Axiom of Revealed Preference (PGARP) if the following two conditions are satisfied.
   a) For every dataset $\mathcal{D}_\tau^i = \{(p_t, x_t^i) : t \in \{1, 2, \ldots, \tau\}\}$ for all $i \in \{1, \ldots, n\}$ and all $\tau \in \{1, \ldots, T\}$, $\mathcal{D}_\tau^i$ satisfies GARP.
   b) The responses x_t^i originated from players in a concave potential game. ■

Note that if only a single agent (i.e. $n = 1$) is considered, then Theorem 4.2 is identical to Afriat's Theorem. Similar to Afriat's Theorem, the constructed concave potential function (41) is ordinal–that is, unique up to positive monotone transformations. Therefore several possible options for $\hat{V}(\cdot)$ exist that would produce identical preference relations to the actual potential function $V(\cdot)$. In 4) of Theorem 4.2, the first condition only provides necessary and sufficient conditions for the dataset $\mathcal{D}$ to be consistent with a Nash equilibrium of a game, therefore the second condition is required to ensure consistency with the other statements in the Multiagent Afriat's Theorem. The intuition that connects statements 1 and 3 in Theorem 4.2 is provided by the

following result from [118]; for any smooth potential game that admits a concave potential function V, a sequence of responses $\{x^i\}_{i\in\{1,2,\ldots,n\}}$ are generated by a pure-strategy Nash equilibrium if and only if it is a maximizer of the potential function,

$$\begin{aligned} & x_t = \{x_t^1, x_t^2, \ldots, x_t^n\} \in \arg\max V(\{x^i\}_{i\in\{1,2,\ldots,n\}}) \\ & \text{s.t.} \quad p_t' x^i \leq I_t^i \qquad \forall i \in \{1,2,\ldots,n\} \end{aligned} \tag{42}$$

for each probe vector $p_t \in \mathbb{R}_+^m$.

The non-parametric test for Nash rationality involves determining if (40) has a feasible solution. Computing parameters v_t and $\lambda_t^i > 0$ in (40) involves solving a linear program with T^2 linear constraints in $(n+1)T$ variables, which has polynomial time complexity [105]. In the special case of one agent, the constraint set in (40) is the dual of the *shortest path problem* in network flows. Using the graph theoretic algorithm presented in [121], the solution of the parameters u_t and λ_t in (35) can be computed with time complexity $O(T^3)$.

D. Learning Algorithm for Response Prediction

In the previous section a non-parametric tests to detect if a dataset $\mathcal{D}$ is consistent with Nash rationality was provided. If the $\mathcal{D}$ satisfies Nash rationality, then the Multiagent Afriat's Theorem can be used to construct the concave potential function of the game for agents in the social network illustrated in Fig.10. In this section we provide a non-parametric learning algorithm that can be used to predict the responses of agents using the constructed concave potential function (41).

To predict the response of agent i, denoted by $\hat{x}_\tau^i$, for probe p_τ and budget I_τ^i, the optimization problem (42) is solved using the estimated potential function $\hat{V}$ (41), p_τ, and I_τ^i. Computing $\hat{x}_\tau^i$ requires solving an optimization problem with linear constraints and concave piecewise linear objective. This can be solved using the interior point algorithm [105]. The algorithm used to predict the response $\hat{x}_\tau = (\hat{x}_\tau^1, \hat{x}_\tau^2, \cdots, \hat{x}_\tau^n)$ is given below:

Step 1: Select a probe vector $p_\tau \in \mathbb{R}_+^m$, and response budget I_τ^i for the estimation of optimal response $\hat{x}_\tau \in \mathbb{R}_+^{m\times n}$.

Step 2: For dataset $\mathcal{D}$, compute the parameters v_t and λ_t^i using (40).

Step 3: The response $\hat{x}_\tau$ is computed by solving the following linear program given $\{\mathcal{D}, p_\tau, I_\tau^i\}$, and $\{v_t, \lambda_t^i\}$ from Step 2:

$$\begin{aligned} \max \quad & z \\ \text{s.t.} \quad & z \leq v_t + \sum_{i=1}^{n} \lambda_t^i p_t'(\hat{x}_\tau^i - x_t^i) \text{ for } t = 1,\ldots,T \\ & p_\tau' \hat{x}_\tau^i \leq I_\tau^i \quad \forall i \in \{1,2,\ldots,n\} \end{aligned} \tag{43}$$

E. Dataset 1: Online Multiwinner Auction

This section illustrates how Afriat's Theorem from Sec.IV-A can be used to determine if bidders in an online multiwinner auction are utility optimizers. Online auctions are rapidly gaining popularity since bidders do not have to gather at the same geographical location. Several researchers have focused on the timing of bids and multiple bidding behavior in Amazon and eBay auctions [122], [123], [124], [125]. The analysis of the bidding behavior can be exploited by auctioneers to target suitable bidders and thereby increase profits.

The multiwinner auction dataset was obtained from an experimental study conducted amongst undergraduate students in Electrical Engineering at Princeton University in March 25$^{\text{th}}$ 2011[19]. The multiwinner auction consists of bidders competing for questions that will aid them for an upcoming midterm exam. The social network is composed of $n = 12$ bidders where the bidders do not interact with other bidders, they only interact with the external influence, refer to Fig.10. Each bidder is endowed with 500 tokens prior to starting the multiwinner auction. The number of questions being auctioned is not known to the bidders, this prevents the bidders from immediately submitting their entire budget in the final auction. Each auction consists of auctioning a single question at an initial price of 10 tokens and has a duration of 30 min. At the beginning of each auction the bidders are provided with the number of winners, denoted k, that auction will have, and the budget of each bidder. The bids are private information with each bidder only informed when their bid has been outbid. The bidders do not communicate with each other during the auction. At the end of each auction, the first k highest bidders are selected, denoted by $\zeta_1, \zeta_2, \ldots, \zeta_k$. The bidders $\zeta_1, \zeta_2, \ldots, \zeta_k$ are awarded with the question, and pay the second largest bid amount (i.e. bidder ζ_k pays ζ_{k-1}'s bid amount). In the multiwinner auction it is in the self interest of bidders to force other bidders to spend too much eliminating them from competing in successive auctions.

If the bidding behaviour of agents satisfies Afriat's test (35) for utility maximization, the next goal is to classify the behaviour of bidders into two categories: strategic and frantic. If a bidder fails Afriat's test then they are classified as irrational. Strategic bidders will typically submit a large number of bids and a smaller bid amount when compared to frantic bidders. With this bidding behaviour, strategic bidders force the other bidders to spend too much eliminating them from competing in future

[19]The experimental data is provided by Leberknight *et al.* [126].

auctions. Frantic bidders are however only interested in winning the current auction.

To apply Afriat's test (35), the external influence p_t, and bidder responses x_t must be defined. The external influence for each bidder is defined by $p_t^i = [p_t^i(1), p_t^i(2)]$ with $p_t^i(1) =$ *initial bid amount* representing the bidders interest level for winning, and $p_t^i(2) =$ *# of winners* representing the perception of winning where i is the bidder and t the auction. Two datasets are considered for analysis denoted by $\mathcal{D}_1$ and $\mathcal{D}_2$. An identical external influence is used to construct both $\mathcal{D}_1$ and $\mathcal{D}_2$. The responses in $\mathcal{D}_1$ are given by $x_t^i = [x_t^i(1), x_t^i(2)]$ where $x_t^i(1) =$ *# of bids* and $x_t^i(2) =$ *mean bid amount*; and for $\mathcal{D}_2$ the inputs of x_t^i are given by $x_t^i(1) =$ *# of bids* and $x_t^i(2) =$ *mean change in bid amount*. The response $x_t^i(2)$ in $\mathcal{D}_1$ provides the expected bid amount, and $x_t^i(2)$ in $\mathcal{D}_2$ a measure of the statistical dispersion of the bids. The budget I_t^i of each bidder has units of *tokens* multiplied by *# of bids*, and is constrained as the number of tokens and auction duration are finite. The datasets $\mathcal{D}_1$ and $\mathcal{D}_2$ are constructed from $T = 6$ auctions. The non-parametric test (35) is applied to each dataset $\mathcal{D}_1$ and $\mathcal{D}_2$ to detect irrational bidders. For dataset $\mathcal{D}_1$ bidder 4 is irrational, and for $\mathcal{D}_2$ bidder 11 is irrational. Note that the classification of irrational behaviour is dependent on the choice of response signals used for analysis by the experimentalist.

For utility maximization bidders, an estimate of the utility function of each bidder is required to classify them as strategic or frantic. To estimate the utility function of the bidders, a subset of data from $\mathcal{D}_1$, denoted as $\bar{\mathcal{D}}_1$, is selected such that the preferences of all agents i in $\bar{\mathcal{D}}_1$ are identical. It was determined that $\bar{\mathcal{D}}_1 = \{(p_t^i, x_t^i) : i \in \{1, 3, 5, 7, 12\}\}$. Since the preferences of these bidders are identical, we can consider all the data in $\bar{\mathcal{D}}_1$ as originating from a single representative bidder. This allows an improved estimate of the utility function of these bidders as compared to learning the utility function for each bidder separately. An analogous explanation is used for the construction of $\bar{\mathcal{D}}_2 = \{(p_t^i, x_t^i) : i \in \{1, 2, 3, 4, 7, 8, 9\}\}$ from the dataset $\mathcal{D}_2$. The estimated utility function for $\bar{\mathcal{D}}_1$ is given in Fig.11(a) and for $\bar{\mathcal{D}}_2$ in Fig.11(b). As seen from Fig.11(a) and Fig.11(b), bidders have a preference to increase the number of bids compared with increasing the mean bid amount or the difference in mean bid amount. This follows logically as $x_t^i(2)$ increases, the bidder will have to pay more tokens to win the question limiting their ability to bid in future auctions. Interestingly, the bidders show strategic and frantic behaviour in both datasets $\bar{\mathcal{D}}_1$ and $\bar{\mathcal{D}}_2$, as seen in Fig.11(a) and Fig.11(b). This is consistent with the results in [124] which show that bidders change their bidding behavior between successive auctions.

The analysis shows that auctioneers should target bidders that show frantic bidding behavior as they are likely to overspend on items increasing the revenue of the auctioneer. Such behavior can be detected using utility maximization test and constructed utility function from Afriat's Theorem.

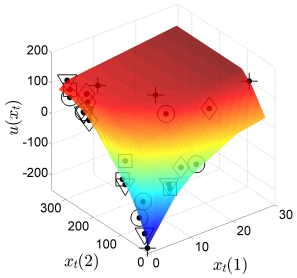

(a) Estimated utility function $u(x_t)$ using dataset $\bar{\mathcal{D}}_1$ defined in Sec.IV-E.

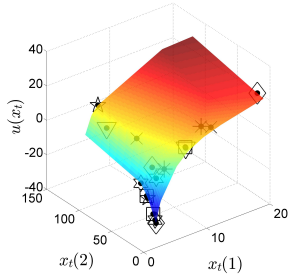

(b) Estimated utility function $u(x_t)$ using dataset $\bar{\mathcal{D}}_2$ defined in Sec.IV-E.

Fig. 11. Estimated utility function of bidders is constructed using (36) using the datasets $\bar{\mathcal{D}}_1$ and $\bar{\mathcal{D}}_2$ defined in Sec.IV-E. The black dots represent the utility associated with experimentally measured responses, and the colour (blue to red) indicates the utility level. The black dots indicate the observed demands, and the shape (i.e. circle, diamond etc.) denotes the respective bidder.

F. Dataset 2: Ontario Electrical Energy Market Dataset

In this section we consider the aggregate power consumption of different zones in the Ontario power grid. A sampling period of $T = 79$ days starting from January 2013 is used to generate the dataset $\mathcal{D}$ for the analysis. All price and power consumption data is available from the *Independent Electricity System Operator*[20] (IESO) website. Each zone is considered as an agent in the corporate network illustrated in Fig.10. The study of corporate social networks was pioneered by Granovetter [127], [128] which shows that the social structure of the network can have important economic outcomes.

[20]http://ieso-public.sharepoint.com/

Examples include agents choice of alliance partners, assumption of rational behavior, self interest behavior, and the learning of other agents behavior. Here we test for rational behavior (i.e. utility maximization and Nash rationality), and if true then learn the associated behavior of the zones. This analysis provides useful information for constructing demand side management (DSM) strategies for controlling power consumption in the electricity market. For example, if a utility function exists it can be used in the DSM strategy presented in [129], [130].

The zones power consumption is regulated by the associated price of electricity set by the senior management officer in each respective zone. Since there is a finite amount of power in the grid, each officer must communicate with other officers in the network to set the price of electricity. Here we utilize the aggregate power consumption from each of the $n = 10$ zones in the Ontario power grid and apply the non-parametric tests for utility maximization (35) and Nash rationality (40) to detect if the zones are demand responsive. If the utility maximization or Nash rationality tests are satisfied, then the power consumption behaviour is modelled by constructing the associated utility function (36) or concave potential function of the game (41).

To perform the analysis the external influence p_t and response of agents x_t must be defined. In the Ontario power grid the wholesale price of electricity is dependent on several factors such as consumer behaviour, weather, and economic conditions. Therefore the external influence is defined as $p_t = [p_t(1), p_t(2)]$ with $p_t(1)$ the average electricity price between midnight and noon, and $p_t(2)$ as the average between noon and midnight with t denoting day. The response of each zone correspond to the total aggregate power consumption in each respective tie associated with $p_t(1)$ and $p_t(2)$ and is given by $x_t^i = [x_t^i(1), x_t^i(2)]$ with $i \in \{1, 2, \ldots, n\}$. The budget I_t^i of each zone has units of dollars as p_t has units of \$/kWh and x_t^i units of kWh.

We found that the aggregate consumption data of each zone does not satisfy Afriat's utility maximization test (35). This points to the possibility that the zones are engaged in a concave potential game–this would not be a surprising result as network congestion games have been shown to reduce peak power demand in distributed demand management schemes [109]. To test if the dataset $\mathcal{D}$ is consistent with Nash rationality the detection test (40) is applied. The dataset for the power consumption in the Ontario power gird is consistent with Nash rationality. Using (43), a concave potential function for the game is constructed. Using the constructed potential function, when do agents prefer to consume power? The *marginal rate of substitution*[21] (MRS) can be used to determine the preferred time for power usage. Formally, the MRS of $x^i(1)$ for $x^i(2)$ is given by

$$\mathrm{MRS}_{12} = \frac{\partial \hat{V}/\partial x^i(1)}{\partial \hat{V}/\partial x^i(2)}.$$

From the constructed potential function we find that $\mathrm{MRS}_{12} > 1$ suggesting that the agents prefer to use power in the time period associated with $x_t(1)$–that is, the agents are willing to give up MRS_{12} kWh of power in the time period associated with $x^i(2)$ for 1 additional kWh of power in time period associated with $x^i(1)$.

The analysis in this section suggests that the power consumption behavior of agents is consistent with players engaged in a concave potential game. Using the Multiagent Afriat's Theorem the agents preference for using power was estimated. This information can be used to improve the DSM strategies presented in [129], [130] to control power consumption in the electricity market.

G. Dataset 3: Twitter Data

Does the tweeting behaviour of Twitter agents satisfy a utility maximization process? The goal is to investigate how tweets and trend indices[22] impact the tweets of agents in the Twitter social network. The information provided by this analysis can be used in social media marketing strategies to improve a brand and for brand awareness. As discussed in [131], Twitter may relay on a huge amount of agent-generated data which can be analyzed to provide novel personal advertising to agents.

To apply Afriat's utility maximization test (35), we choose the external influence and response as follows. External influence $p_t = [\textit{\#Sony}, 1/\textit{\#Playstation}]$ for each day t. The associated response taken by the agents in the network is given by $x_t = [\textit{\#Microsoft}, \textit{\#Xbox}]$. Notice that the probe $p_t(2)$ can be interpreted as the frequency of tweets with the word *Playstation* (i.e. the trending index). The dataset $\mathcal{D}$ of external influence and responses is constructed from $T = 80$ days of Twitter data starting from January 1st 2013. The dataset $\mathcal{D}$ satisfies the utility maximization test (35). This establishes that utility function exists for agents that is dependent on the number of tweets containing the words *Microsoft* and *Xbox*. The data shows that tweets containing the word *Microsoft* and *Xbox* are dependent on the number of tweets containing *Sony* and trending index of *Playstation*. This dependency

[21]The amount of one good that an agent is willing to give up in exchange for another good while maintaining the same level of utility.

[22]Here we define the *trend index* as the frequency of tweets containing a particular word [131].

is expected as Microsoft produces the game console Xbox, and Sony produces the game console Playstation both which have a large number of brand followers (e.g. Xbox has over 3 million, and Playstation over 4 million). To gain further insight into the behaviour of the agents, (36) from Afriat's Theorem is used to construct a utility function for the agents. Fig.12 shows the constructed utility function of the agents. As seen, agents have a higher utility for using the word *Microsoft* as compared to *Xbox*–that is, agents prefer to use the word *Microsoft* to that of *Xbox*. Interestingly, if we define the response to be $x_t = [\#Microsoft, 1/\#Xbox]$, then the dataset satisfies utility maximization. From the constructed utility function, not shown, the agents prefer to increase the tweets containing the word *Microsoft* compared to increasing the trend index of *Xbox*. If instead $x_t = [1/\#Microsoft, 1/\#Xbox]$, then the dataset satisfies utility maximization and agents prefer to increase the trend index of *Microsoft* compared to that of *Xbox*.

To summarize, the above analysis suggests the following interesting fact: *Xbox* has a lower utility than *Microsoft* in terms of Twitter sentiment. Therefore, online marketing strategies should target the brandname *Microsoft* instead of *Xbox*.

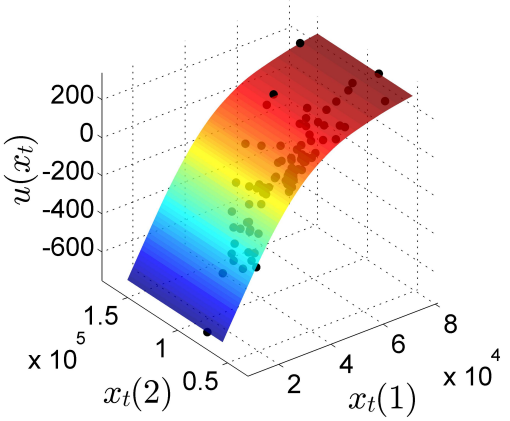

Fig. 12. Estimated utility function $u(x_t)$ using dataset $\mathcal{D}$ defined in Sec.IV-G and constructed using the non-parametric learning algorithm (36) from Afriat's Theorem.

H. Summary and Extensions

The principle of revealed preferences is an active research area with numerous recent papers. We have already mentioned the papers [102], [30], [103]. Below we summarize some related literature that extends the basic framework of Afriat's theorem.

Afriat's theorem holds for finite datasets and gives an explicit construction of a class of concave utility functions that rationalize the dataset. Mas Colell [132] has given sufficient conditions under which as the data set size T grows to infinity, the underlying utility function of the consumer can be fully identified.

Though the classical Afriat's theorem holds for linear budget constraints $p_t' x \leq I_t$ in (34), an identical formulation holds for certain non-linear budget constraints as illustrated in [133]. The budget constraints considered in [133] are of the form $\{x \in \mathbb{R}^m_+ | g(x) \leq 0\}$ where $g : \mathbb{R}^m_+ \rightarrow \mathbb{R}$ is an increasing continuous function and $\mathbb{R}^m_+$ denotes the positive orthant. Also [133] shows how the results in [132] on recoverability of the utility function can be extended to such nonlinear budget constraints. However, learning the utility function from a finite dataset in the case of a non-linear budget constraint requires sophisticated machine learning algorithms [134]. The machine learning algorithms can only guarantee that the estimated utility function is approximately consistent with the dataset $\mathcal{D}$–that is, the estimated utility is not guaranteed to contain all the preference relations consistent with the dataset $\mathcal{D}$.

In [121], results in statistical learning theory are applied to the principle of revealed preferences to address the question: Is the class of demand functions derived from monotone concave utilities efficiently probably approximately correct (PAC) learnable? It is shown that Lipschitz utility functions are efficiently PAC learnable. In [120], the authors extend the results of [121] and show that for agents engaged in a concave potential game that satisfy Nash rationality, if the underlying potential function satisfies the Lipschitz condition then the potential function of the game is PAC learnable.

Finally, in many cases, the responses of agents are observed in noise. Then determining if an agent is a utility maximizer (or a multiagent system's response is consistent with play from a Nash) becomes a statistical decision test. In [135] it is shown how stochastic optimization algorithms can be devised to optimize the probe signals to minimize the type II errors of the decision test subject to a fixed type I error.

V. Closing Remarks

This paper has discussed three important and interrelated themes regarding the dynamics of social sensors, namely, diffusion models for information in social networks, Bayesian social learning and revealed preferences. In each case, examples involving real datasets were given to illustrate the various concepts. The unifying theme behind these three topics stems from predicting global behavior given local behavior: individual

social sensors make decisions and learn from other social sensors and we are interested in understanding the behavior of the entire network. In Sec.II we showed that the global degree of infected nodes can be determined by mean field dynamics. In Sec.III, it was shown that despite the apparent simplicity in information flows between social sensors, the global system can exhibit unusual behavior such as herding and data incest. Finally, in Sec.IV a non-parametric method was used to determine the utility functions of a multiagent system - this can be used to predict the response of the system.

This paper has dealt with social sensing issues of relevance to an engineering audience. There are several topics of relevance to social sensors that are omitted due to space constraints, including:

- Coordination of decisions via game-theoretic learning [136], [137], [138] or Bayesian game models such as global games [139]
- Consensus formation over social networks and cooperative models of network formation [140], [141]
- Small world models [142], [143]
- Peer to peer media sharing [144], [145]
- Privacy and security modelling [146], [147]

REFERENCES

[1] A. Rosi, M. Mamei, F. Zambonelli, S. Dobson, G. Stevenson, and J. Ye, "Social sensors and pervasive services: Approaches and perspectives," in *Proceedings of the 2011 IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOM Workshops)*. IEEE, 2011, pp. 525–530.

[2] R. Lee and K. Sumiya, "Measuring geographical regularities of crowd behaviors for Twitter-based geo-social event detection," in *Proceedings of the 2nd ACM SIGSPATIAL International Workshop on Location Based Social Networks*. ACM, 2010, pp. 1–10.

[3] Z. Cheng, J. Caverlee, and K. Lee, "You are where you tweet: A content-based approach to geo-locating Twitter users," in *Proceedings of the 19th ACM International Conference on Information and Knowledge Management*. ACM, 2010, pp. 759–768.

[4] M. Trusov, A. V. Bodapati, and R. E. Bucklin, "Determining influentual users in internet social networks," *Journal of Marketing Research*, vol. XLVII, pp. 643–658, Aug. 2010.

[5] J. Leskovec, L. A. Adamic, and B. A. Huberman, "The dynamics of viral marketing," *ACM Transactions on the Web (TWEB)*, vol. 1, no. 1, p. 5, 2007.

[6] T. Sakaki, M. Okazaki, and Y. Matsuo, "Earthquake shakes Twitter users: real-time event detection by social sensors," in *Proceedings of the 19th International Conference on World Wide Web*. New York, NY, USA: ACM, 2010, pp. 851–860.

[7] J. Bollen, H. Mao, and X. Zeng, "Twitter mood predicts the stock market," *Journal of Computational Science*, vol. 2, no. 1, pp. 1–8, 2011.

[8] B. Pang and L. Lee, "Opinion mining and sentiment analysis," *Foundations and Trends in Information Retrieval*, vol. 2, no. 1-2, pp. 1–135, 2008.

[9] S. Asur and B. A. Huberman, "Predicting the future with social media," in *2010 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT)*, vol. 1. IEEE, 2010, pp. 492–499.

[10] B. Ifrach, C. Maglaras, and M. Scarsini, "Monopoly pricing in the presence of social learning," *NET Institute Working Paper No. 12-01*, 2011.

[11] M. Luca, *Reviews, reputation, and revenue: The case of Yelp.com, Technical Report 12-016*. Harvard Business School, September 2011.

[12] S. Goel, D. Watts, and D. Goldstein, "The structure of online diffusion networks," in *Proceedings of the 13th ACM conference on Electronic Commerce*. ACM, 2012, pp. 623–638.

[13] N. Bennett. (2011, April) 4 great (free) tools to measure social sentiment and 4 important stats. [Online]. Available: http://socialmediatoday.com/nick-bennett/287405/4-great-free-tools-measure-social-sentiment-and-4-important-stats

[14] E. Sun, I. Rosenn, C. Marlow, and T. M. Lento, "Gesundheit! modeling contagion through facebook news feed." in *Proceedings of the 3rd International AAAI Conference on Weblogs and Social Media*, San Jose, CA, May 2009.

[15] D. M. Romero, B. Meeder, and J. Kleinberg, "Differences in the mechanics of information diffusion across topics: Idioms, political hashtags, and complex contagion on Twitter," in *Proceedings of the 20th International Conference on World Wide Web*, Hyderabad, India, Mar. 2011, pp. 695–704.

[16] C. Chamley, *Rational herds: Economic Models of Social Learning*. Cambridge University Press, 2004.

[17] D. López-Pintado, "Diffusion in complex social networks," *Games and Economic Behavior*, vol. 62, no. 2, pp. 573–590, 2008.

[18] R. Pastor-Satorras and A. Vespignani, "Epidemic spreading in scale-free networks," *Physical Review Letters*, vol. 86, no. 14, p. 3200, 2001.

[19] M. Benaïm and J. Weibull, "Deterministic approximation of stochastic evolution in games," *Econometrica*, vol. 71, no. 3, pp. 873–903, 2003.

[20] A. Banerjee, "A simple model of herd behavior," *Quaterly Journal of Economics*, vol. 107, no. 3, pp. 797–817, August 1992.

[21] S. Bikchandani, D. Hirshleifer, and I. Welch, "A theory of fads, fashion, custom, and cultural change as information cascades," *Journal of Political Economy*, vol. 100, no. 5, pp. 992–1026, October 1992.

[22] D. Acemoglu and A. Ozdaglar, "Opinion dynamics and learning in social networks," *Dynamic Games and Applications*, vol. 1, no. 1, pp. 3–49, 2011.

[23] I. Lobel, D. Acemoglu, M. Dahleh, and A. Ozdaglar, "Preliminary results on social learning with partial observations," in *Proceedings of the 2nd International Conference on Performance Evaluation Methodolgies and Tools*. Nantes, France: ACM, 2007.

[24] D. Acemoglu, M. Dahleh, I. Lobel, and A. Ozdaglar, "Bayesian learning in social networks," National Bureau of Economic Research, Working Paper 14040, May 2008.

[25] T. Cover and M. Hellman, "The two-armed-bandit problem with time-invariant finite memory," *IEEE Transactions on Information Theory*, vol. 16, no. 2, pp. 185–195, 1970.

[26] M. Hellman and T. Cover, "Learning with finite memory," *The Annals of Mathematical Statistics*, vol. 41, no. 3, pp. 765–782, 1970.

[27] C. Chamley, A. Scaglione, and L. Li, "Models for the diffusion of beliefs in social networks: An overview," *IEEE Signal Processing Magazine*, vol. 30, no. 3, pp. 16–29, 2013.

[28] V. Krishnamurthy and H. V. Poor, "Social learning and Bayesian games in multiagent signal processing: How do local and global decision makers interact?" *IEEE Signal Processing Magazine*, vol. 30, no. 3, pp. 43–57, 2013.

[29] S. Afriat, "The construction of utility functions from expenditure data," *International economic review*, vol. 8, no. 1, pp. 67–77, 1967.

[30] H. Varian, "The nonparametric approach to demand analysis," *Econometrica*, vol. 50, no. 1, pp. 945–973, 1982.

[31] ——, "Price discrimination and social welfare," *The American Economic Review*, pp. 870–875, 1985.

[32] ——, "Revealed preference and its applications," *The Economic Journal*, vol. 122, no. 560, pp. 332–338, 2012.

[33] ——, "Online ad auctions," *The American Economic Review*, pp. 430–434, 2009.

[34] S. Haykin, "Cognitive radio: Brain-empowered wireless communications," *IEEE Journal on Selected Areas Communications*, vol. 23, no. 2, pp. 201–220, Feb. 2005.

[35] M. D. Choudhury, S. Counts, and M. Czerwinski, "Identifying relevant social media content: leveraging information diversity and user cognition," in *Proceedings of the 22nd ACM conference on Hypertext and hypermedia*. ACM, 2011, pp. 161–170.

[36] J. Kleinberg, "Cascading behavior in networks: Algorithmic and economic issues," *Algorithmic game theory*, vol. 24, pp. 613–632, 2007.

[37] F. Vega-Redondo, *Complex Social Networks*. Cambridge University Press, 2007, vol. 44.

[38] M. Jackson, *Social and Economic Networks*. Princeton University Press, 2010.

[39] D. Easley and J. Kleinberg, *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge University Press, 2010.

[40] K. C. Chen, M. Chuang, and H. V. Poor, "From technological networks to social networks," *IEEE Journal on Selected Areas Communications*, vol. 31, no. 9, pp. 548–572, Sept. 2013.

[41] V. Krishnamurthy and H. V. Poor, "A tutorial on interactive sensing in social networks," *IEEE Transactions on Computational Social Systems*, vol. 1, no. 1, pp. 3–21, March 2014.

[42] V. Krishnamurthy, O. N. Gharehshiran, and M. Hamdi, "Interactive sensing and decision making in social networks," *Foundations and Trends® in Signal Processing*, vol. 7, no. 1-2, pp. 1–196, 2014.

[43] M. Granovetter, "Threshold models of collective behavior," *American Journal of Sociology*, vol. 83, no. 6, pp. 1420–1443, May 1978.

[44] E. Mossel and S. Roch, "On the submodularity of influence in social networks," in *Proceedings of 39th Annual ACM Symposium on Theory of Computing*. ACM, 2007, pp. 128–134.

[45] N. Chen, "On the approximability of influence in social networks," *SIAM Journal on Discrete Mathematics*, vol. 23, no. 3, pp. 1400–1415, 2009.

[46] D. López-Pintado, "Contagion and coordination in random networks," *International Journal of Game Theory*, vol. 34, no. 3, pp. 371–381, 2006.

[47] D. Aldous, "Interacting particle systems as stochastic social dynamics," *Bernoulli*, vol. 19, no. 4, pp. 1122–1149, 2013.

[48] J. Leskovec, "SNAP library," http://snap.stanford.edu/data/index.html.

[49] P. Erdös and A. Rényi, "On random graphs, I," *Publicationes Mathematicae (Debrecen)*, vol. 6, pp. 290–297, 1959.

[50] A. Barabasi and A. Reka, "Emergence of scaling in random networks," *Science*, vol. 286, no. 5439, p. 509, 1999.

[51] R. Elliott, L. Aggoun, and J. Moore, *Hidden Markov Models – Estimation and Control*. New York: Springer-Verlag, 1995.

[52] A. Benveniste, M. Metivier, and P. Priouret, *Adaptive Algorithms and Stochastic Approximations*, ser. Applications of Mathematics. Springer-Verlag, 1990, vol. 22.

[53] H. Kushner and G. Yin, *Stochastic Approximation Algorithms and Recursive Algorithms and Applications*, 2nd ed. Springer-Verlag, 2003.

[54] G. Yin, V. Krishnamurthy, and C. Ion, "Regime switching stochastic approximation algorithms with application to adaptive discrete stochastic optimization," *SIAM Journal on Optimization*, vol. 14, no. 4, pp. 117–1215, 2004.

[55] A. Culotta, "Lightweight methods to estimate influenza rates and alcohol sales volume from twitter messages," *Language resources and evaluation*, vol. 47, no. 1, pp. 217–238, 2013.

[56] A. Signorini, A. Segre, and P. Polgreen, "The use of twitter to track levels of disease activity and public concern in the us during the influenza a h1n1 pandemic," *PloS one*, vol. 6, no. 5, p. e19467, 2011.

[57] F. Zhang, J. Luo, C. Li, X. Wang, and Z. Zhao, "Detecting and analyzing influenza epidemics with social media in china," in *Advances in Knowledge Discovery and Data Mining*. Springer, 2014, pp. 90–101.

[58] A. Culotta, "Towards detecting influenza epidemics by analyzing twitter messages," in *Proceedings of the first workshop on social media analytics*. ACM, 2010, pp. 115–122.

[59] ——, "Lightweight methods to estimate influenza rates and alcohol sales volume from twitter messages," *Language resources and evaluation*, vol. 47, no. 1, pp. 217–238, 2013.

[60] J. Li and C. Cardie, "Early stage influenza detection from twitter," *arXiv preprint arXiv:1309.7340*, 2013.

[61] D. Broniatowski, M. Paul, and M. Dredze, "National and local influenza surveillance through twitter: An analysis of the 2012-2013 influenza epidemic," *PloS one*, vol. 8, no. 12, p. e83672, 2013.

[62] H. Achrekar, A. Gandhe, R. Lazarus, S. Yu, and B. Liu, "Twitter improves seasonal influenza prediction," in *HEALTHINF 2012 International Conference on Health Informatics*, 2012, pp. 61–70.

[63] ——, "Predicting flu trends using twitter data," in *2011 IEEE Conference on Computer Communications Workshops*. IEEE, 2011, pp. 702–707.

[64] N. Christakis and J. Fowler, "Social network sensors for early detection of contagious outbreaks," *PloS one*, vol. 5, no. 9, p. e12948, 2010.

[65] J. Leskovec and C. Faloutsos, "Sampling from large graphs," in *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. Philadelphia: ACM press, 2006, pp. 631–636.

[66] A. Dasgupta, R. Kumar, and D. Sivakumar, "Social sampling," in *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*. Beijing: ACM, 2012, pp. 235–243.

[67] D. D. Heckathorn, "Respondent-driven sampling: a new approach to the study of hidden populations," *Social Problems*, vol. 44, pp. 174–199, 1997.

[68] ——, "Respondent-driven sampling ii: deriving valid population estimates from chain-referral samples of hidden populations," *Social Problems*, vol. 49, pp. 11–34, 2002.

[69] S. Lee, "Understanding respondent driven sampling from a total survey error perspective," *Survey Practice*, vol. 2, no. 6, 2009.

[70] M. Malekinejad, L. Johnston, C. Kendall, L. Kerr, M. Rifkin, and G. Rutherford, "Using respondent-driven sampling methodology for HIV biological and behavioral surveillance in international settings: a systematic review," *AIDS and Behavior*, vol. 12, no. S1, pp. 105–130, 2008.

[71] S. Goel and M. J. Salganik, "Respondent-driven sampling as Markov chain Monte Carlo," *Statistics in Medicine*, vol. 28, pp. 2209–2229, 2009.

[72] A. Lansky, A. Abdul-Quader, M. Cribbin, T. Hall, T. Finlayson, R. Garffin, L. S. Lin, and P. Sullivan, "Developing an HIV behavioral surveillance system for injecting drug users: the National HIV Behavioral Surveillance System," *Public Health Reports*, vol. 122, no. S1, pp. 48–55, 2007.

[73] J. Kleinberg and P. Raghavan, "Query incentive networks," in *46th Annual IEEE Symposium onFoundations of Computer Science, 2005. FOCS 2005.* IEEE, 2005, pp. 132–141.

[74] A. Doucet, N. Gordon, and V. Krishnamurthy, "Particle filters for state estimation of jump Markov linear systems," *IEEE Transactions on Signal Processing*, vol. 49, pp. 613–624, 2001.

[75] A. Logothetis and V. Krishnamurthy, "Expectation maximization algorithms for MAP estimation of jump Markov linear systems," *IEEE Transactions on Signal Processing*, vol. 47, no. 8, pp. 2139–2156, August 1999.

[76] V. Krishnamurthy and M. Hamdi, "Mis-information removal in social networks: Dynamic constrained estimation on directed acyclic graphs," *IEEE Journal Selected Topics in Signal Processing*, vol. 7, no. 2, pp. 333–346, May 2013.

[77] B. Tuttle, "Fact-checking the crowds: How to get the most out of hotel-review sites," *Time Magazine*, July 2013.

[78] G. Adomavicius and A. Tuzhilin, "Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions," *IEEE Transactions on Knowledge and Data Engineering*, vol. 17, no. 6, pp. 734–749, 2005.

[79] I. Konstas, V. Stathopoulos, and J. M. Jose, "On social networks and collaborative recommendation," in *Proceedings of the 32nd International ACM SIGIR conference on Research and Development in Information Retrieval*. ACM, 2009, pp. 195–202.

[80] Y. Kanoria and O. Tamuz, "Tractable Bayesian social learning on trees," in *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, Jul. 2012, pp. 2721 –2725.

[81] J. Pearl, "Fusion, propagation, and structuring in belief networks," *Artificial Intelligence*, vol. 29, no. 3, pp. 241 – 288, 1986.

[82] K. Murphy, Y. Weiss, and M. Jordan, "Loopy belief propagation for approximate inference: an empirical study," in *Proceedings of the Fifteenth Conference Uncertainty in Artificial Intelligence*, 1999, pp. 467–475.

[83] J. Yedidia, W. Freeman, and Y. Weiss, "Constructing free-energy approximations and generalized belief propagation algorithms," *IEEE Transactions on Information Theory*, vol. 51, no. 7, pp. 2282–2312, 2005.

[84] R. J. Aumann, "Agreeing to disagree," *The Annals of Statistics*, vol. 4, no. 6, pp. 1236–1239, Nov. 1976.

[85] J. Geanakoplos and H. Polemarchakis, "We can't disagree forever," *Journal of Economic Theory*, vol. 28, no. 1, pp. 192 – 200, 1982.

[86] V. Borkar and P. Varaiya, "Asymptotic agreement in distributed estimation," *IEEE Transactions on Automatic Control*, vol. 27, no. 3, pp. 650–655, 1982.

[87] R. Bond, C. Fariss, J. Jones, A. Kramer, C. Marlow, J. Settle, and J. Fowler, "A 61-million-person experiment in social influence and political mobilization," *Nature*, vol. 489, pp. 295–298, September 2012.

[88] D. Kempe, J. Kleinberg, and E. Tardos, "Maximizing the spread of influence through a social network," in *Proceedings of the 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Washington, DC, Aug 2003, pp. 137–146.

[89] V. Krishnamurthy, "Quickest detection POMDPs with social learning: Interaction of local and global decision makers," *IEEE Transactions on Information Theory*, vol. 58, no. 8, pp. 5563–5587, 2012.

[90] S. Karlin, *Total Positivity*. Stanford Univ., 1968, vol. 1.

[91] S. Karlin and Y. Rinott, "Classes of orderings of measures and related correlation inequalities. I. Multivariate totally positive distributions," *Journal of Multivariate Analysis*, vol. 10, no. 4, pp. 467–498, December 1980.

[92] P. Milgrom, "Good news and bad news: Representation theorems and applications," *Bell Journal of Economics*, vol. 12, no. 2, pp. 380–391, 1981.

[93] D. Topkis, *Supermodularity and Complementarity*. Princeton University Press, 1998.

[94] P. Milgrom and C. Shannon, "Monotone comparative statistics," *Econometrica*, vol. 62, no. 1, pp. 157–180, 1992.

[95] R. Amir, "Supermodularity and complementarity in economics: An elementary survey," *Southern Economic Journal*, vol. 71, no. 3, pp. 636–660, 2005.

[96] M. Hamdi and V. Krishnamurthy, "Removal of Data Incest in Multi-agent Social Learning in Social Networks," *ArXiv e-prints*, Sep. 2013.

[97] J. Surowiecki, *The Wisdom of Crowds*. New York: Anchor, 2005.

[98] C. Anderson and G. J. Kilduff, "Why do dominant personalities attain influence in face-to-face groups? the competence-signaling effects of trait dominance," *Journal of Personality and Social Psychology*, vol. 96, no. 2, pp. 491–503, 2009.

[99] M. Ottaviani and P. Sørensen, "Information aggregation in debate: Who should speak first?" *Journal of Public Economics*, vol. 81, no. 3, pp. 393–421, 2001.

[100] P. Samuelson, "A note on the pure theory of consumer's behaviour," *Economica*, pp. 61–71, 1938.

[101] S. Afriat, *Logic of choice and economic theory*. Clarendon Press Oxford, 1987.

[102] W. Diewert, "Afriat and revealed preference theory," *The Review of Economic Studies*, pp. 419–425, 1973.

[103] R. Blundell, "How revealing is revealed preference?" *Journal of the European Economic Association*, vol. 3, no. 2-3, pp. 211–235, 2005.

[104] A. Fostel, H. Scarf, and M. Todd, "Two new proofs of Afriats theorem," *Economic Theory*, vol. 24, no. 1, pp. 211–219, 2004.

[105] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.

[106] H. Varian, "Non-parametric tests of consumer behaviour," *The Review of Economic Studies*, vol. 50, no. 1, pp. 99–110, 1983.

[107] R. Deb, "Interdependent preferences, potential games and household consumption," University Library of Munich, Germany, MPRA Paper 6818, Jan 2008.

[108] A. Chapman, G. Verbic, and D. Hill, "A healthy dose of reality for game-theoretic approaches to residential demand response," in *2013 IREP Symposium Bulk Power System Dynamics and Control-IX Optimization, Security and Control of the Emerging Power Grid (IREP)*. IEEE, 2013, pp. 1–13.

[109] C. Ibars, M. Navarro, and L. Giupponi, "Distributed demand management in smart grid with a congestion game," in *2010 First IEEE International Conference on Smart Grid Communications*. IEEE, 2010, pp. 495–500.

[110] C. Wu, H. Mohsenian-Rad, J. Huang, and A. Wang, "Demand side management for wind power integration in microgrid using dynamic potential game theory," in *2011 IEEE GLOBECOM Workshops*, 2011, pp. 1199–1204.

[111] R. Rosenthal, "A class of games possessing pure-strategy Nash equilibria," *International Journal of Game Theory*, vol. 2, no. 1, pp. 65–67, 1973.

[112] A. Hayrapetyan, E. Tardos, and T. Wexler, "The effect of collusion in congestion games," in *Proceedings of the Thirty-eighth Annual ACM Symposium on Theory of Computing*, ser. STOC '06. ACM, 2006, pp. 89–98.

[113] M. Babaioff, R. Kleinberg, and C. Papadimitriou, "Congestion games with malicious players," in *Proceedings of the 8th ACM Conference on Electronic Commerce*, ser. EC '07. New York, NY, USA: ACM, 2007, pp. 103–112.

[114] P. McGill and M. Schumaker, "Boundary conditions for single-ion diffusion," *Biophysical Journal*, vol. 71, pp. 1723–1742, 1996.

[115] H. Varian, "Revealed preference," *Samuelsonian economics and the twenty-first century*, pp. 99–115, 2006.

[116] R. Deb, "A testable model of consumption with externalities," *Journal of Economic Theory*, vol. 144, no. 4, pp. 1804 – 1816, 2009.

[117] T. Ui, "Correlated equilibrium and concave games," *International Journal of Game Theory*, vol. 37, pp. 1–13, 2008.

[118] A. Neyman, "Correlated equilibrium and potential games," *Int. J. Game Theory*, vol. 26, no. 2, pp. 223–227, 1997.
[119] A. Carvajal, R. Deb, J. Fenske, and J. Quah, "Revealed preference tests of the cournot model," *Econometrica*, vol. 81, no. 6, pp. 2351–2379, 2013.
[120] W. Hoiles and V. Krishnamurthy, "Nonparametric demand forecasting and detection of demand-responsive consumers," *IEEE Transactions on Smart Grid*.
[121] E. Beigman and R. Vohra, "Learning from revealed preference," in *Proceedings of the 7th ACM Conference on Electronic Commerce*, ser. EC '06. New York, NY, USA: ACM, 2006, pp. 36–42.
[122] S. Borle, P. Boatwright, and J. Kadane, "The timing of bid placement and extent of multiple bidding: An empirical investigation using ebay online auctions," *Statistical Science*, pp. 194–205, 2006.
[123] A. Ockenfels and A. Roth, "The timing of bids in internet auctions: Market design, bidder behavior, and artificial agents," *AI magazine*, vol. 23, no. 3, p. 79, 2002.
[124] R. Pownall and L. Wolk, "Bidding behavior and experience in internet auctions," *European Economic Review*, vol. 61, pp. 14–27, 2013.
[125] S. DellaVigna, "Psychology and economics: Evidence from the field," *Journal of Economic Literature*, vol. 47, no. 2, pp. 315–72, 2009.
[126] C. Leberknight, H. Inaltekin, M. Chiang, and H. Poor, "The evolution of online social networks: A tutorial survey," *IEEE Signal Processing Magazine*, vol. 29, no. 2, pp. 41–52, 2012.
[127] M. Granovetter, "Economic action and social structure: the problem of embeddedness," *American journal of sociology*, pp. 481–510, 1985.
[128] ——, "The impact of social structure on economic outcomes," *Journal of economic perspectives*, pp. 33–50, 2005.
[129] S. Dave, M. Sooriyabandara, and L. Zhang, "Application of a game-theoretic energy management algorithm in a hybrid predictive-adaptive scenario," in *2011 2nd IEEE PES International Conference and Exhibition on Innovative Smart Grid Technologies (ISGT Europe)*, 2011, pp. 1–6.
[130] H. Nguyen, J. Song, and Z. Han, "Demand side management to reduce peak-to-average ratio using game theory in smart grid," in *2012 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, 2012, pp. 91 –96.
[131] F. Giummolè, S. Orlando, and G. Tolomei, "A study on microblog and search engine user behaviors: How twitter trending topics help predict google hot queries," *HUMAN*, vol. 2, no. 3, p. 195, 2013.
[132] A. Mas-Colell, "On revealed preference analysis," *The Review of Economic Studies*, pp. 121–131, 1978.
[133] F. Forges and E. Minelli, "Afriat's theorem for general budget sets," *Journal of Economic Theory*, vol. 144, no. 1, pp. 135–145, 2009.
[134] S. Lahaie, "Kernel methods for revealed preference analysis," in *ECAI 2010: 19th European Conference on Artificial Intelligence*, 2010, pp. 439–444.
[135] V. Krishnamurthy and W. Hoiles, "Afriat's test for detecting malicious agents," *IEEE Signal Processing Letters*, vol. 19, no. 12, pp. 801–804, 2012.
[136] S. Hart and A. Mas-Colell, "A simple adaptive procedure leading to correlated equilibrium," *Econometrica*, vol. 68, no. 5, pp. 1127–1150, 2000.
[137] S. Hart, A. Mas-Colell, and Y. Babichenko, *Simple Adaptive Strategies: From Regret-Matching to Uncoupled Dynamics*, ser. World Scientific Series in Economic Theory. World Scientific Publishing, 2013, vol. 4.
[138] O. Namvar, V. Krishnamurthy, and G. Yin, "Distributed tracking of correlated equilibria in regime switching noncooperative games," *IEEE Transactions on Automatic Control*, vol. 58, no. 10, pp. 2435–2450, 2013.
[139] G. Angeletos, C. Hellwig, and A. Pavan, "Dynamic global games of regime change: Learning, multiplicity, and the timing of attacks," *Econometrica*, vol. 75, no. 3, pp. 711–756, 2007.
[140] B. Kozma and A. Barrat, "Consensus formation on adaptive networks," *Physical Review E*, vol. 77, no. 1, p. 016102, 2008.
[141] A. Tahbaz-Salehi and A. Jadbabaie, "Consensus over ergodic stationary graph processes," *IEEE Transactions on Automatic Control*, vol. 55, no. 1, pp. 225–230, Jan. 2010.
[142] D. J. Watts, "Networks, dynamics, and the small-world phenomenon," *American Journal of Sociology*, vol. 105, no. 2, pp. 493–527, Sep. 1999.
[143] J. Kleinberg, "Navigation in a small world," *Nature*, vol. 406, no. 6798, pp. 845–845, Aug. 2000.
[144] M. J. Halvey and M. T. Keane, "Exploring social dynamics in online media sharing," in *Proceedings of 16th international conference on World Wide Web*, Banff, AB, Canada, May 2007, pp. 1273–1274.
[145] H. V. Zhao, W. S. Lin, and K. J. R. Liu, *Behavior dynamics in media-sharing social networks*. Cambridge University Press, 2011.
[146] P. G. Lange, "Publicly private and privately public: Social networking on youtube," *Journal of Computer-Mediated Communication*, vol. 13, no. 1, pp. 361–380, Oct. 2007.
[147] L. Liu, E. Yu, and J. Mylopoulos, "Security and privacy requirements analysis within a social setting," in *Proceedings of 11th IEEE International Requirements Engineering Conference*, Monterey Bay, CA, Sep. 2003, pp. 151–161.

Vikram Krishnamurthy [F] (vikramk@ece.ubc.ca) is a professor and Canada Research Chair at the Department of Electrical Engineering, University of British Columbia, Vancouver, Canada. Dr Krishnamurthy's current research interests include statistical signal processing, computational game theory and stochastic control in social networks. He served as distinguished lecturer for the IEEE Signal Processing Society and Editor in Chief of IEEE Journal Selected Topics in Signal Processing. He received an honorary doctorate from KTH (Royal Institute of Technology), Sweden in 2013.

William Hoiles (whoiles@ece.ubc.ca) is currently a Ph.D. candidate in the Department of Electrical and Computer Engineering, University of British Columbia, Vancouver, Canada. He received the M.A.Sc. degree in 2012 from the Department of Engineering Science, Simon Fraser University, Vancouver, Canada. His current research interests are social sensors and the bioelectronic interface.